

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representation of
The original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

This Page Blank (uspto)



10526 U.S. PTO

09/575193



Patent Office
Canberra

I, LEANNE MYNOTT, TEAM LEADER EXAMINATION SUPPORT AND SALES hereby certify that annexed is a true copy of the Provisional specification in connection with Application No. PQ0559 for a patent by SILVERBROOK RESEARCH PTY LTD filed on 25 May 1999.

WITNESS my hand this
Eleventh day of April 2000

LEANNE MYNOTT
TEAM LEADER EXAMINATION
SUPPORT AND SALES



This Page Blank (uspto)

AUSTRALIA
Patents Act 1990

PROVISIONAL SPECIFICATION

Applicant(s) :

SILVERBROOK RESEARCH PTY LTD

Invention Title:

A METHOD AND APPARATUS (NPAGE01)

The invention is described in the following statement:

A METHOD AND APPARATUS (NPAGE01)

Field of the Invention

The present invention relates to the field of information distribution and printing.

Background of the Invention

5 The background of the invention is detailed in the attached appendix A which also sets out a detailed description of the invention.

Summary of the Invention

 It is an object of the present invention to provide an improved information distribution system.

10 In accordance with a first aspect of the present invention, there is provided an information distribution system for distributing customized information on print media, the system comprising a first series of data collection units for collecting customizable data; a network interconnecting the data collection units with a series of customization output printers; a series of customization output printers comprising: a user identification means to identify a user of the customization output printer; a printer
15 activation means to activate the customization output printer; a pagewidth printer element for printing customized booklets; such that, upon activation of the user identification means and the printer activation means, the data collection units download a current customized booklet for the user and the printer element prints out the customized booklet.

 The customization output printer can be arranged such that pages of the booklet are preferably
20 fed through the printer from a top to bottom with the booklet being output at a bottom of the printer. The customization output printers can include a hopper located in the bottom thereof and the booklet can be gravity dropped into the hopper. The customization output printer further preferably can include a binding means for binding the pages of the booklet together. The binding means can comprise a fluid ejection device ejecting a glue onto a predetermined portion of substantially each page of the booklet.
25 The binding means further can comprise compression means for compressing each page of the booklet together.

 The user identification means can comprise a fingerprint scanner. The pagewidth printer can comprise an elongated wall mounted unit. The data collection units utilize multicasting facilities in downloading data to the customization output printers which select stories from the data collection units
30 depending upon user interest. The customization output printers utilize public key cryptography for access control with the data collection units.

 In accordance with a further aspect of the present invention, there is provided an information distribution system for distributing customized information on print media, the system comprising a first series of data collection units for collecting customizable data; a network interconnecting the data
35 collection units with a series of customization output printers; a series of customization output printers comprising: a user identification means to identify a user of the customization output printer; a printer activation means to activate the customization output printer; a pagewidth printer element for printing

customized booklets; booklet collation means for collating the booklet; a print roll unit for storing rolls of print media for utilization for printing by the pagewidth printer element; such that, upon activation of the user identification means and the printer activation means, the data collection units download a current customized booklet for the user and the printer element prints out on the print roll unit pages of the
5 customized booklet which are formed into a booklet by the booklet collation means.

In accordance with a further aspect of the present invention, there is provided an interactive information distribution system comprising a series of interconnected data collection units for collecting information; an interconnected network; a series of printers for printing out information; a pen device including an image sensor for sensing an image area; wherein the image printed preferably can include a
10 series of visually distinctive areas containing identification information such that, upon placing the pen in the vicinity of a visually distinctive area, the pen is activated to sense and decode the identification information, the identification information is forwarded to the distribution computers which in turn creates a further customized page for forwarding to a printer unit for printing on demand of a user.

The identification information can be stored in conjunction with a user identification for
15 determining likely future news items of interest to the user. The information preferably can include advertising customized in accordance with a users selection of particular identifiers.

In accordance with a further aspect of the present invention, there is provided a news distribution system comprising a series of automated news collection units for collecting, collating and storing news stories; a distribution network for electronic distribution of news stories to a series of news
20 printers; a series of news printers interconnected to the news collection units via the distribution network for printing out a collection of news stories on demand on print media; the printing out including the printing out of a series of detectable location markers on the print media; a handheld scanner unit adapted to scan the detectable location markers to determine a current scanning location on the print media, the handheld scanner being interconnected to the distribution network for transmission of the current
25 scanning locations to the automated news collection units; the automated news collection units utilizing the current scanning locations to determine customized future news stories to send to a user of the news printer.

The detectable location markers are preferably substantially invisible to the human eye. The detectable location markers can comprise a regularly spaced grid in conjunction with a unique
30 identifiable positional marker.

The news stories are preferably page formatted by the automated news collection units and downloaded in a compressed form over the distribution network. The news stories are preferably composited together into a page by the news printers.

The output printers are preferably double sided and included watermark data on images printed
35 out, the watermark data being detected by means of interference between a front and back surface of the print media.

In accordance with a further aspect of the present invention, there is provided a news distribution system comprising a series of automated news collection units for collecting, collating and storing news stories; a distribution network for electronic distribution of news stories to a series of news printers; a series of news printers interconnected to the news collection units via the distribution network for printing out a collection of news stories on demand on print media; the printing out including the
5 printing out of a series of detectable location markers on the print media; a handheld scanner unit adapted to scan the detectable location markers to determine a current scanning location on the print media, the handheld scanner being interconnected to the distribution network for transmission of the current scanning locations to the automated news collection units; the automated news collection units utilizing
10 the current scanning locations to determine customized future news stories to send to a user of the news printer. The handheld scanner unit can comprise a pen shaped unit.

The system can further comprise a charging unit for placing the pen shaped unit in to recharge the pen shaped unit. The charging unit can be cup shaped. The charging unit preferably can include a communications unit adapted to register pens placed in the charging unit with the printer. The
15 handheld scanner unit preferably can include a series of accelerometers to determine a path taken by the handheld scanner unit. The handheld scanner unit preferably can include signature recognition means for recognizing a signature pattern executed by a user of the scanner unit. The pen shaped unit preferably can include an activation light which changes color in accordance with the operation being carried out by the pen shaped unit.

20

Description Preferred in Other Embodiments

The preferred embodiment of the present invention is as set out in the attached appendix A which provides for a detailed description of the implementation of an information distribution and printing system denoted NETPAGE.

25 The ink jet technology utilized can be a wide pagewidth printing system as disclosed in the attached appendix B.

Appendix C illustrates an example customized personal newspaper printed out by the NETPAGE system.

30 Appendix D illustrates an example of NETPAGE interaction to order a series of hardware items.

Appendix E illustrates an example of NETPAGE interaction to conduct an email interaction.

35 It would be appreciated by a person skilled in the art that numerous variations and/or modifications may be made to the present invention as shown in the specific embodiments in the attached documentation without departing from the spirit or scope of the invention as broadly described in the attached appendices. The present embodiments are, therefore, to be considered in all respects to be illustrative and not restrictive.

We Claim:

1. An information distribution system for distributing customized information on print media, said system comprising:
 - a first series of data collection units for collecting customizable data;
 - 5 a network interconnecting said data collection units with a series of customization output printers;
 - a series of customization output printers comprising:
 - a user identification means to identify a user of said customization output printer;
 - a printer activation means to activate said customization output printer;
 - 10 a pagewidth printer element for printing customized booklets;
 - such that, upon activation of said user identification means and said printer activation means, said data collection units download a current customized booklet for said user and said printer element prints out said customized booklet.
2. A system as claimed in claim 1 wherein said customization output printer is
15 arranged such that pages of said booklet are fed through said printer from a top to bottom with the booklet being output at a bottom of said printer.
3. A system as claimed in claim 2 wherein said customization output printers include a hopper located in the bottom thereof and said booklet is gravity dropped into said hopper.
4. A system as claimed in claim 1 wherein said customization output printer
20 further includes a binding means for binding the pages of said booklet together.
5. A system as claimed in claim 4 wherein said binding means comprises a fluid ejection device ejecting a glue onto a predetermined portion of substantially each page of said booklet.
6. A system as claimed in claim 5 wherein said binding means further comprises compression means for compressing each page of said booklet together.
- 25 7. A system as claimed in claim 1 wherein said user identification means comprises a fingerprint scanner.
8. A system as claimed in claim 1 wherein said pagewidth printer comprises an elongated wall mounted unit.
9. A system as claimed in any previous claim wherein said data collection units
30 utilize multicasting facilities in downloading data to said customization output printers which select stories from said data collection units depending upon user interest.
10. A system as claimed in claim 1 wherein said customization output printers utilize public key cryptography for access control with said data collection units.
11. An information distribution system for distributing customized information on
35 print media, said system comprising:
 - a first series of data collection units for collecting customizable data;
 - a network interconnecting said data collection units with a series of customization

output printers;

a series of customization output printers comprising:

a user identification means to identify a user of said customization output printer;

a printer activation means to activate said customization output printer;

5 a pagewidth printer element for printing customized booklets;

booklet collation means for collating said booklet;

a print roll unit for storing rolls of print media for utilization for printing by said
pagewidth printer element;

10 such that, upon activation of said user identification means and said printer activation
means, said data collection units download a current customized booklet for said user and said printer
element prints out on said print roll unit pages of said customized booklet which are formed into a
booklet by said booklet collation means.

12. An interactive information distribution system comprising:

a series of interconnected data collection units for collecting information;

15 an interconnected network;

a series of printers for printing out information;

a pen device including an image sensor for sensing an image area;

20 wherein the image printed includes a series of visually distinctive areas containing
identification information such that, upon placing the pen in the vicinity of a visually distinctive area,
the pen is activated to sense and decode the identification information, the identification information is
forwarded to said distribution computers which in turn creates a further customized page for forwarding to
a printer unit for printing on demand of a user.

25 13. A system as claimed in claim 12 wherein said identification information is
stored in conjunction with a user identification for determining likely future news items of interest to said
user.

14. A system as claimed in claim 13 wherein said information includes advertising
customized in accordance with a users selection of particular identifiers.

15. A news distribution system comprising:

30 a series of automated news collection units for collecting, collating and storing news
stories;

a distribution network for electronic distribution of news stories to a series of news
printers;

35 a series of news printers interconnected to said news collection units via said
distribution network for printing out a collection of news stories on demand on print media;
said printing out including the printing out of a series of detectable location markers on
said print media;

a handheld scanner unit adapted to scan said detectable location markers to determine a

current scanning location on said print media, said handheld scanner being interconnected to said distribution network for transmission of said current scanning locations to said automated news collection units;

5 said automated news collection units utilizing said current scanning locations to determine customized future news stories to send to a user of said news printer.

16. A system as claimed in claim 15 wherein said detectable location markers are substantially invisible to the human eye.

17. A system as claimed in claim 15 or 16 wherein said detectable location markers comprise a regularly spaced grid in conjunction with a unique identifiable positional marker.

10 18. A system as claimed in claim 15 wherein said news stories are page formatted by said automated news collection units and downloaded in a compressed form over said distribution network.

19. A system as claimed in claim 15 wherein said news stories are composited together into a page by said news printers.

15 20. A system as claimed in claim 15 wherein said output printers are double sided and included watermark data on images printed out, said watermark data being detected by means of interference between a front and back surface of the print media.

21. A news distribution system comprising:
a series of automated news collection units for collecting, collating and storing news stories;
20 a distribution network for electronic distribution of news stories to a series of news printers;

a series of news printers interconnected to said news collection units via said distribution network for printing out a collection of news stories on demand on print media;
25 said printing out including the printing out of a series of detectable location markers on said print media;

a handheld scanner unit adapted to scan said detectable location markers to determine a current scanning location on said print media, said handheld scanner being interconnected to said distribution network for transmission of said current scanning locations to said automated news collection units;
30

 said automated news collection units utilizing said current scanning locations to determine customized future news stories to send to a user of said news printer.

22. A system as claimed in claim 21 wherein said handheld scanner unit comprises a pen shaped unit.

35 23. A system as claimed in claim 22 further comprising:
a charging unit for placing said pen shaped unit in to recharge said pen shaped unit.

24. A system as claimed in claim 23 wherein said charging unit is cup shaped.

25. A system as claimed in claim wherein said charging unit includes a communications unit adapted to register pens placed in said charging unit with said printer.

26. A system as claimed in claim 21 wherein said handheld scanner unit includes a series of accelerometers to determine a path taken by said handheld scanner unit.

5 27. A system as claimed in claim 21 wherein said handheld scanner unit includes signature recognition means for recognizing a signature pattern executed by a user of said scanner unit.

28. A system as claimed in claim 22 wherein said pen shaped unit includes an activation light which changes color in accordance with the operation being carried out by said pen shaped unit.

10 29. A network news distribution system substantially as hereinbefore described with reference to the attached appendices.

30. A network news distribution system as claimed in any previous claim when utilizing an ink jet print head having a series of spaced apart thermal bend actuators for the ejection of ink.


ABSTRACT

An information distribution system for distributing customized information on print media, the system comprising a first series of data collection units for collecting customizable data; a network
5 interconnecting the data collection units with a series of customization output printers; a series of customization output printers comprising: a user identification means to identify a user of the customization output printer; a printer activation means to activate the customization output printer; a pagewidth printer element for printing customized booklets; such that, upon activation of the user
10 identification means and the printer activation means, the data collection units download a current customized booklet for the user and the printer element prints out the customized booklet.

This Page Blank (uspto)

Netpage System Design Description

preliminary version 0.1, 23 May 1999



Silverbrook Research Pty Ltd
393 Darling Street, Balmain
NSW 2041 Australia
Phone: +61 2 9818 6633
Fax: +61 2 9818 6711
Email: info@silverbrook.com.au

Appendix A

Document History

Version	Date	Authors	Details
0.1p	23 May 1999	Paul Lapstun Toby King Simon Walmsley Kia Silverbrook	Preliminary issue.

Contents

OVERVIEW	1
1 Introduction	2
1.1 Publishing to Netpage Printers	3
2 The Demise of Paper	5
3 News and Advertising Trends	7
3.1 Newspapers	7
3.2 Online News Delivery	8
3.3 Online Advertising	9
3.4 Online Classified Advertising	10
4 News Personalization	11
ARCHITECTURE	14
5 Netpage System Architecture	15
5.1 The Internet	15
5.2 Netpages and Netpage Documents	16
5.3 The Netpage Network	17
5.4 The Netpage Printer	18
5.5 The Netpage Pen	20
5.6 Netpage Interaction	21
5.7 Standard Features of Netpages	23
5.8 The Help System	23
6 Personalized Publication Model	25
6.1 Editorial Personalization	25
6.2 Advertising Localization and Targeting	26
6.3 User Profiles	28
6.4 Intelligent Page Layout	29
6.5 Document Format	29
6.6 Document Distribution	30
7 Security	32
7.1 Cryptography	32
7.2 Netpage Printer Security	33
7.3 Netpage Pen Security	33
7.4 Secure Documents	34
7.5 Non-Repudiation	35
8 Electronic Commerce Model	36
8.1 Secure Electronic Transaction (SET)	36
8.2 SET Payments	36
8.3 Micro-Payments	37
APPLICATIONS AND BUSINESS MODELS	38
9 Applications	39
9.1 Personalized Subscriptions	39
9.2 On-Demand Publications	40
9.3 E-Commerce	41

9.4	Communication.....	41
9.5	Web Browsing and Printing	42
9.6	Corporate Intranets.....	42
9.7	Government Services	43
9.8	Photo Album.....	43
9.9	Computer Printing.....	43
10	Business Models.....	44
10.1	System Principles	44
10.2	Bootstrapping	45
10.3	Maturity.....	46
	NETPAGE PRINTER.....	48
11	Printer Product Design.....	49
11.1	Wallprinter.....	49
11.2	Wallprinter Pro	54
11.3	Tableprinter Pro	57
11.4	Wallprinter Pro R	58
11.5	Tableprinter Pro R.....	61
11.6	Travelprinter.....	62
11.7	Netpage Pen.....	64
11.8	Pen Charging Cup	67
12	Memjet-Based Printing	68
13	Document Data Flow.....	69
14	Printer Controller Architecture	71
14.1	Detailed Document Data Flow	72
14.2	Print Engine Architecture.....	73
15	Pen Controller Architecture	75
	DETAILED DESIGN	76
16	Position Tag Data Format.....	77
16.1	Position Tag Structure.....	79
16.2	Position Tag Encoding	82
17	In-Pen Processing.....	85
17.1	Position Tag Decoding.....	85
17.2	Data Encryption	98
	REFERENCES	99

OVERVIEW

1 Introduction

Netpages are pages of high-quality text, graphics and images printed on ordinary paper, but which work almost like an interactive Web pages. Information encoded on each page in invisible ink is picked up by an optically-imaging pen and transmitted to the network. Active “links” and “buttons” on each page can be “pressed” with the pen to request information from the network or signal preferences to a server. Text written by hand on a Netpage is automatically recognized via the pen, allowing forms to be filled in. Signatures recorded on a Netpage are automatically verified, allowing e-commerce transactions to be securely authorized.

The pen works in conjunction with a Netpage Printer, an Internet-wired printing appliance for the home or office. The pen is wireless and communicates with the Netpage Printer using an encrypted radio frequency signal.

The Netpage Printer delivers, periodically or on demand, personalized newspapers, magazines, brochures and other publications, all printed at high quality on interactive Netpages. Unlike a personal computer, the Netpage Printer is an appliance typically wall-mounted in the kitchen or near the breakfast table, i.e. the place where the morning news is first consumed, and the household’s point of departure for the day. It also comes in a low-power, wireless, portable model.

Netpages printed at their point of consumption combine the ease-of-use of paper with the timeliness and interactivity of an interactive medium.

Netpages are crucially enabled by Memjet printing technology, which makes high-speed magazine-quality printing affordable to consumers. A Netpage publication has the physical characteristics of a traditional news magazine, i.e. a set of letter-size glossy pages printed in full color on both sides, bound together for easy navigation and comfortable handling. The Netpage Printer prints 60 to 90 full-color Netpages per minute.

The Netpage Printer exploits the growing availability of broadband Internet access. Cable service is available to 95% of households in the United States [87], and cable modem service offering broadband Internet access is already available to 20% of these [8]. The Netpage Printer also operates with slower connections, but with longer delivery times and lower image quality.

Netpage Publication Servers on the Internet deliver print-quality publications to Netpage Printers. Periodical publications are delivered automatically to subscribing Netpage Printers via pointcasting and multicasting Internet protocols. Personalized publications are filtered and formatted according to individual user profiles.

A Netpage Printer supports any number of pens, and a pen can work with any number of Netpage Printers. Each pen has a unique identifier. By default, a Netpage Pen identifies a particular user mostly by convention. A household may have a collection of colored Netpage Pens, one assigned to each member of the family. This allows each user to maintain a distinct profile with respect to a Netpage Publication Server. A Netpage Pen can also be registered with an Identity Server and linked to one or more payment card accounts. This allows e-commerce payments to be securely authorized using the Netpage Pen. The Identity Server compares the signature captured by the Netpage Pen with a previously registered signature, allowing it to authenticate the user’s identity to the e-commerce system. Other biometrics can also be used to verify identity. A version of the Netpage Pen includes fingerprint scanning, verified in a similar way by the Identity Server.

Although a Netpage Printer delivers periodicals such as the morning newspaper without user intervention, it never delivers unsolicited junk mail. It only delivers periodicals from subscribed or otherwise authorized sources. The Netpage Printer is unlike a fax machine or e-mail account which is visible to any junk mailer who knows the telephone number or email address.

1.1 PUBLISHING TO NETPAGE PRINTERS

There are a number of advantages to publishing to Netpage Printers. The magazine quality of Netpage Printer output makes it a more attractive publishing and advertising medium than both traditional newsprint and computer screens.

The cost of paper and ink consumption is transferred to the user. Subscription fees can be eliminated entirely in lieu of the user taking on this extra cost, and the user thus perceives a better-value product. The erratic price of newsprint is removed from publishers' profit equations, resulting in more stable margins. A new market for paper and ink consumables, with its own margins, is created.

The cost of consumables can be selectively subsidized, for example when non-editorial publications such as product brochures and account statements are printed.

Capital and maintenance expenditure on printing plant is effectively transferred to the user, although the perceived expense is small because Netpage Printers are sold at close to cost or given away to encourage adoption, subsidized by future advertising profits. Maintenance can also be subsidized or its cost included in a longer term service agreement.

Costly physical distribution is replaced by electronic distribution via a preexisting and widely subscribed network - the Internet.

Both the editorial and advertising content of publications delivered via the Netpage Network can be customized for each user. Editorial content can be personalized according to the user's profile. Advertising can be localized to the user's locality and optionally targeted to the user's demographic.

A personalized publication can be a small fraction of the size of its traditionally-delivered counterpart, yet contain the same amount of information relevant to the user, and in a more accessible form. The user appreciates the more efficient and digestible publication.

Localized advertising can be targeted to more specific localities and their associated demographics, and this allows advertising space to be exploited more efficiently, i.e. with less waste. Advertisers are constantly pressing traditional publishers for greater localization, something which they have great difficulty delivering cost-effectively.

By revealing personal information such as age, gender, marital status, income, profession, education, etc., the user can allow the advertising to be more carefully targeted. In return they can receive greater subsidies and discounted product offers. As advertising becomes more targeted, it becomes less of a nuisance and more of a service in itself.

Although the publication's per-page circulation figures fall drastically, the publication's actual per-section readership is preserved, and the correspondingly higher advertising rates for personalized delivery can exactly compensate for this.

Advertising delivered via the Netpage Network has the dual benefits of print and online delivery. Print supports the impact of large-format ads. Online delivery supports customization, linking, and measurability, and consequently online charging models.

Consider a full-page advertisement for a new car model in a news magazine delivered via the Netpage Network. The advertising campaign can be national or even international. The ad only appears if compatible with the user's demographic, either implied by their ZIP code or more explicitly by their personal details. Anyone who requests a product brochure via the on-ad button receives one immediately via their Netpage Printer, customized with a list of local dealers. If they press a particular dealer's "contact me" button in the brochure, the dealer receives a message via the system and contacts the user by telephone.

The publisher profits in the normal way by selling the advertising space, but can also profit by receiving a fee on the "click-through" to the brochure, and a commission on any product sale which eventuates.

The Netpage Network promises to be the most effective advertising medium ever conceived. It combines the editorial and print quality of traditional publications with arbitrarily finely targeted advertising, and provides a direct link between advertising, product information, and purchasing. Added revenue from click-through fees and e-commerce commissions may even allow users' costs - printer, ink, paper, and Internet access - to be fully subsidized.

2 The Demise of Paper

Online publication has many advantages over traditional paper-based publication. From the consumer's point of view, information is available on demand, information can be navigated via hypertext links, information can be searched, and information can be automatically personalized.

From the publisher's point of view, the costs of printing and physical distribution are eliminated, and the publication becomes more attractive to the advertisers who pay for it because it can be targeted to specific demographics and linked to product sites.

Online publication also has a few disadvantages. Computer screens are inferior to paper. At the same quality as a magazine page, an SVGA computer screen displays only about a fifth as much information¹. Both CRTs and LCDs have brightness and contrast problems, particularly when ambient light is strong. Ink on paper, being reflective rather than emissive, is both bright and sharp in ambient light.

Faced with reading more than the most trivial amounts of text on a screen, most people prefer to print it before reading it. Increasingly, online publishers are recognizing this and providing information in formats suitable for printing. At one extreme this means providing text-only versions of documents so they print efficiently, i.e. without imposing a screen format on the printed page; at the other extreme it means providing formatted versions of documents - e.g. in Adobe's Portable Document Format (PDF) - so they print at high quality.

Editorial content is often compromised to fit the online medium and the habits of the people who frequent it, who tend to browse rather than read. Although powerful new advertising models become possible, it becomes more difficult to deliver effective advertisements.

To truly enable online publication, many people envisage a universal information appliance - a lightweight portable "tablet" with a page-size touch-sensitive color display and a high-bandwidth wireless connection to the Internet. First proposed by Xerox PARC's Alan Kay in the mid 1970s in the form of the "Dynabook", and partially realized in recent paperback-sized electronic books, even Bill Gates is now confidently predicting that such a device will soon augur the death of paper publications [28].

To achieve low power consumption, low weight, and paper-like display quality, a bistable reflective display technology is required. Several candidates are now emerging from the labs, including Kent Display's cholesteric LCD technology (chLCD) [47], Xerox' "Gyricon" rotating ball technology [36], and E Ink's electrophoretic technology [24,67]. ChLCD is arguably closest to practical deployment [23].

Next-generation cellular phone networks promise 2Mbps packet switching [25], comparable to the broadband access people are getting used to in cable networks [63]. Satellite networks, while offering or promising still higher speeds [39,72,93], require receivers difficult to deploy in mobile devices.

Beyond the vision of the basic tablet, E Ink imagines its digital ink "printed" onto a number of flexible pages bound into a book, preserving the physical navigability of a paper-based publication, and approaching its low cost, but allowing the pages to be rewrit-

1. Assuming a magazine page has an equivalent digital resolution of 200 continuous-tone pixels per inch.

ten electronically in place. They optimistically predict newspapers delivered in this way within five years [17], despite fundamental problems yet to be overcome [23].

The advantages of a tablet are many. Unlike a desktop or notebook computer, a tablet may actually provide a pleasant reading experience. Unlike a paper publication, a tablet provides intelligent access to an unlimited amount of information; its weight is not dictated by the amount of information it carries. More than just an information appliance, it can also act as a multi-purpose multimedia communications device and interactive entertainment device.

A tablet has disadvantages too. It uses batteries which run down and have to be recharged. It may break when dropped or malfunction when exposed to hot coffee. It's not quite cheap enough to be disposable - so there's still a problem if it's misplaced or stolen. It has a "user interface" which has to be learned. The leading candidate display technology - chLCD - is still less than half as reflective (i.e. "bright") as paper.

The drawbacks of traditional paper-based publications have little to do with paper itself, and much to do with how the information gets onto the printed page. The economics of centralized printing and distribution prevent the kind of information selection, personalization and navigation people have come to expect from interactive electronic media such as the Internet. The inefficiency of printing and distributing a hundred-page newspaper to a customer who may read only a few pages is widely decried.

Given a technology such as Memjet, it becomes economic to print high-quality publications at their point of consumption rather than at their point of production. The Netpage Printer leverages Memjet to deliver personalized publications to the home, gaining many of the advantages of online publication, while retaining the ease-of-use of high-quality printed ink on paper.

Netpages and the Netpage Printer address the key problems of online publication, without relying on the development and consumer acceptance of a new reading device.

3 News and Advertising Trends

3.1 NEWSPAPERS

People obtain news from a variety of sources - network and cable television, radio, daily newspapers, and weekly newsmagazines. In the United States, although the various news media are healthy and profitable, per capita news consumption is somewhat in decline as a new generation of young adults have less time to read and favor television entertainment over news [18]. Yet six out of ten adult Americans read a newspaper every day [77].

The United States has about 1500 daily newspapers with a total circulation of 57 million. Just the top ten "national" dailies (Wall Street Journal, USA Today, New York Times, Los Angeles Times, Washington Post, etc.) account for a circulation of 10 million. The major weekly newsmagazines (Time, Newsweek, U.S. News) have a similar (weekly) circulation of about 10 million.

In 1997, newspaper companies' revenue exceeded \$24 billion, a five-year high, and margins nudged 20% [52], due both to increased spending on advertising and to reduced prices of newsprint.

Television and radio, by their nature, excel at delivering breaking news. Newspapers and newsmagazines, on the other hand, deliver the depth and analysis behind the headlines. Broadcast news in isolation does a poor job of informing the public. The more local the news is, the poorer the broadcast coverage, and the greater the public's dependence on newspapers.

Newspaper content and packaging has evolved considerably since the 1970s. News is somewhat softer, news stories are shorter and more well-written, there are more feature articles, and there is more editorial and reader opinion. Newspapers are more structured. Identifiable sections make them more accessible, and provide greater focus for advertisers. Much special-interest content has migrated from daily inclusion to weekly sections. These cover topics such as lifestyle, personal finance, entertainment, technology, etc. The proportion of graphics and pictures is greater. Color is widely used. Newspapers are easier to use and more entertaining than ever before, if at the expense of some "hard" news.

Daily newspapers are growing increasingly dependent on the various wire services. A newspaper may excel at local and regional news, but rely on the major wires (Associated Press and United Press International) for national and international news, the so-called "supplemental" wires (LA Times/Washington Post, NY Times, Scripps-Howard, Knight-Ridder Tribune, etc.) for specific strengths (and value-for-money), and the international wires (Reuters etc.) for international perspective. A growing number of newspapers operate more as news aggregators than news gatherers.

Advertising typically contributes more than 75% of newspaper and magazine revenue, while subscriptions contribute less than 25% [31,73,74]. National advertising makes up roughly 14% of advertising spending, retail advertising 46%, and classified advertising 40% [74].

Advertisers are pursuing increasingly specific targeting, favoring quality newspaper readership over raw circulation [27], and using more targeted media where possible. Magazines, for example, have more specific readerships than newspapers, free "shoppers" are

very localized by their nature, while direct mailers can target demographics based on individually-categorized ZIP codes, or databases of individuals.

Newspapers have responded with geographically zoned editions to support local advertising, and greater sectioning of their product. They have also expanded their page counts to provide more advertising scope, despite erratic newsprint prices in the 1990s [3].

Despite this, there is ongoing conflict between newspapers' mass distribution model, and advertisers' need for micro-targeting [30]. This conflict, coupled with advertisers' desire for higher-quality printing of color images, is motivating a shift from run-of-press (ROP) advertising to inserts [74].

3.2 ONLINE NEWS DELIVERY

Fearing the online migration of advertising, traditional news publishers from both broadcast and print have ventured into Internet-based news delivery, wanting to establish a presence at whatever cost before newcomers become entrenched. Most newspapers are still reporting losses from their online operations [74].

Online news delivery offers a number of advantages. Breaking news can be delivered as soon as it happens. News can be customized for individual readers according to their preferences and geographic locations. Readers can explore stories to arbitrary depth, follow links to related resources, and search archival material. Readers can participate in discussion groups and contribute to opinion polls. The news itself can incorporate audio and video clips, and can include live transmissions, converging with broadcast.

Online news delivery also has disadvantages. Computer screens are of limited size and quality compared with print. Few people enjoy reading a story of any length on a computer screen. Computers are not portable in the wide sense that a newspaper is. The news may be more timely, but the time and place in which it can be consumed are more constrained than with a newspaper.

Despite the power of hypertext, many online readers express a preference for a linear presentation, "where they [can] skim one section after another until the presentation [is] exhausted" [10]. Interestingly, a majority of traditional newspaper readers admit they scan every page in the main section of the newspaper [77], looking for items of interest without necessarily knowing what they're looking for, and achieving some kind of closure at the end. Online hypertext, by contrast, is both a limitless resource and a bottomless pit.

While traditional news publishers such as The New York Times can deploy full editorial content online [94], newcomers such as Yahoo typically only provide "raw" news items sourced from the wire services [70].

A recent survey indicates that 21% of the 74 million Internet users in the United States regularly read news online as an alternative to traditional print and broadcast sources, and 16% obtain a major proportion of their news online [78]. More broadly, between 37% and 64% of the Internet population reads news online at least once a week. The fluctuations in the figures are related to what may be happening in the news. Major or breaking news stories attract more users - 46% of Americans say they only follow national news stories when "something major is happening" [26].

With 41% of Americans online, the Internet population has become mainstream, and the weather has become the most popular news online. This is closely followed by technology

news, entertainment news, and local news. As one observer puts it, all of this “sound[s] like the 6 o’clock news” [78]. As a reflection of these habits, the online audience share of national newspapers has diminished from 23% in 1995 to 16% in 1998, while the online audience share of broadcast TV sites has grown.

3.3 ONLINE ADVERTISING

At its simplest, advertising alerts a motivated customer to the availability of a product, possibly at a competitive price. At a more sophisticated level, advertising seeks to influence future purchasing decisions by creating brand awareness. Ultimately, advertising seeks to create desire for a product even when actual need is absent.

Advertising prices are traditionally based on how many people see the advertisement, and their spending power in relation to the product. In practice, the more homogeneous the demographics of the audience, the easier it is to match to a product, and hence the higher the corresponding advertising cost per thousand (CPM). Broadcast media use ratings and timeslot demographics to set advertising rates. Print media use audited circulation figures and sectional readership demographics.

The simplest online advertising model is also based on how many people see the ad. Online this has the advantage of being based on solid numbers, since the number of “impressions” of a particular Web page can be counted exactly.

The specific advantage of an online ad, however, is that the ad itself can measurably capture a sales lead by acting as a link to a product site. The product site may simply provide more product information in the form of specifications, pricing, and ordering details. It may also support immediate online ordering, thus completing the link from ad to sale. Beyond providing simple ad exposure, it is this measurable linking of advertisement to sales lead or sale which is the strength of online advertising [84]. Cost per click (CLC) charging is gaining acceptance but is still controversial.

Beyond CLC, there exists the possibility of paying a commission to the ad host on any sale that actually eventuates [88]. Amazon.com is probably the best-known example of a company paying commissions to other sites in this way.

The broader advantage of online advertising is that advertising can be localized and targeted arbitrarily finely, in conjunction with the publication of online content such as news. This is the strategy pursued by online advertising agencies such as Click-Through [15], which acts as the middle-man between advertisers and online content publishers. They expect online advertising to represent more than 10 percent of all advertising revenue by 2001.

Since online ads are necessarily small-format, they communicate best with motivated customers already on the look-out for a particular product or service. Online ads are less suited to building brand awareness or creating buying desires, since the real substance of the advertisement - the product Web site - is a click away from the initial small-format ad. A small-format online ad can’t provide the single-hit emotional impact of a large-format print ad, and conversely, the online world can’t support the large-format ads that print can.

So-called interstitial ads, which appear full-screen when traversing from one page of information to the next, go some way to providing a medium for larger-format ads online [88]. User resistance, however, seems to be preventing their widespread use.

3.4 ONLINE CLASSIFIED ADVERTISING

Classified advertising is indisputably suited to online delivery. Unlike their traditional printed counterparts, online classifieds can be easily searched, and are not subject to space constraints. The online migration of classified advertising is considered a serious threat to newspapers' classified advertising profits [90], and some newspapers are building an online presence for this reason alone. Some observers predict as much as 50% of classified advertising revenue moving online within the next ten years [106].

Another problem faced by newspapers, who rely on classifieds for up to 40% of advertising spending, is that many newcomers are offering free online classified advertising as a way of building a venue for non-classified advertising.

4 News Personalization

From the reader's point of view, a personalized news publication can provide more information in fewer pages. The actual form this personalization takes, however, is not necessarily obvious.

The MIT Media Lab's News in the Future (NiF) project has been championing the concept of "The Daily Me" for almost two decades [65]. Nicholas Negroponte, one of the project's founders, envisages a highly personalized news publication which is no longer driven by "what other people think is news" [71]. By way of examples close to his own needs, it includes news about people and places about to be encountered, and puts "the most important [news] of all" - a summary of e-mail - on the front page [6]. Negroponte recognizes the need to vary the degree of personalization, advocating a higher "serendipity factor" on a lazy Sunday than on a working weekday.

The opposing view holds that the value of a news publication lies precisely in its *shared* nature. It reflects the common concerns and values of a community of readers, and establishes a baseline of expectations of what they are all supposed to know [97]. As a consequence, the publication also speaks with a consistent editorial voice and with consistent assumptions about the reader's level of background knowledge. Such a shared publication allows its readers to orient themselves in relation to their community.

NiF's Walter Bender answers the charge (in his own words) of "The Daily Me engendering a fragmented world populated by self-interested myopes", by stressing the possibilities of personalizing individual news items [4]. This can consist in varying the depth of an item, or supplementing it with background information, based on the reader's level of knowledge. It can involve interpreting information relative to the reader's background, such as (somewhat dubiously) making value judgements about the weather relative to the reader's normal home town weather. It can also be as simple as using metric rather than imperial units.

FishWrap [10,66], MIT's personalized campus newspaper and NiF's latest offering, goes further by creating a front page whose content represents an explicit community consensus. Each front page news item is prioritized according to the number of readers who put it forward for inclusion. The rest of the newspaper is still personalized according to each reader's profile, consisting of reader-defined sections containing topics of interest.

There are two implications of recognizing the shared nature of news. Firstly, some news is news to everybody in a community, no matter how personalized they claim they would like their news to be. This implies that the community must make decisions about news item priority, either directly (as in FishWrap) or indirectly via a proxy (i.e. an editor).

Secondly, a news item can only be properly understood in the context of the community for which it is intended. This implies that a news item must be branded with its source (assuming that the source implies the intended target). As an example, it is significant whether a news item regarding the proof of Fermat's last theorem is branded with New Scientist or The New York Times. To a professional mathematician, the latter implies, by its very existence, that the proof is of significance beyond the scientific community.

Of course, a news item must also be branded to allow its source to build and maintain its brand. The brand then allows the reader to infer the quality of the news item from the known quality of the source.

Most personalization of news uses *feature-based filtering*. This means that news item content is matched to topics and keywords in the reader's profile. News sources tag the items they produce with various information to allow them to be effectively filtered. This tagging may be brief or extensive, and may include such things as news item urgency, byline, news category, subject(s), keyword(s), date and time, and location [43]. The body text of a news item can also be scanned directly for keywords, but this may result in false matches if keywords are interpreted out of context. Items in the text such as personal names and locations can be tagged to reduce such ambiguity [43]. Similarly, dates, times, and monetary amounts can be tagged to allow localized presentation.

Feature-based filtering suffers from a number of problems. Filtering based on tags is only as good as the original tagging. The latest tagging standards are only just beginning to be adopted [7]. Filtering based on the text itself is constrained by the intelligence of the text parsing. If based simply on keyword matching, it can be both inaccurate, generating false matches because of word sense ambiguity, and imprecise, generating false mismatches because of a lack of inference.

Feature-based filtering is incapable of discerning more abstract attributes such as quality, style, and point-of-view (unless they're indicated by tags). And since it only matches items anticipated by the user's profile, it is a poor generator of serendipitous finds.

FishWrap's front page comes into existence based on a crude form of *collaborative filtering*. In its broader form, collaborative filtering involves sharing recommendations (or ratings) among *like-minded* people [86]. This means that one person's ratings influence another person if and only if the two share similar interests, i.e. they have similar rating histories. Collaborative filtering overcomes many of the problems of feature-based filtering, since ratings originate with people who have digested the items in question, rather than from automated analysis of the items. Collaborative filtering sidesteps the issue of *explaining* why a person might like a particular item.

Collaborative filtering has problems of its own. The system only works if people are willing to contribute ratings. In contributing ratings, of course, they are both doing the community a service and tuning their own interest profiles. The statistical error in correlating people's interests decreases as the number of ratings increases. However, incentives may have to be offered to encourage people to contribute ratings.

To bootstrap the accumulation of ratings for new items, an independent mechanism must exist to distribute them to a critical mass of people. Conversely, to bootstrap the accumulation of interest profiles for new users, an independent mechanism must exist to distribute a critical mass of items to them.

To allow meaningful accumulation of ratings, a sufficient period of time must be allowed to elapse. This may conflict with the timely delivery of items in question.

The statistical correlation between different people's interests, represented by their rating histories, is most meaningful when the ratings apply to homogeneous items. For a set of heterogeneous items, collaborative filtering is best applied to homogeneous subsets.

In a news setting, collaborative filtering is best applied to feature articles. Features have the longer life span required to support the accumulation of ratings, and are often appreciated for abstract qualities best singled out by collaborative filtering (good writing, humor, incisiveness, etc.).

Naturally, the larger a publication's readership, and the better its taste in relation to a potential reader, the stronger the publication's brand will appear to that reader.

Although it's easy to become preoccupied with automatic filtering, in reality there's more to editing the news than just filtering news feeds. An editor also solicits news, commissions analysis, and offers opinion, ideally ensuring that the publication offers a balanced and complete view of the world.

Perhaps the most important personalization step a reader takes is in selecting a particular publication from a set of available publications, based on its perceived quality and relevance.

Thus the publication's brand equates to the highest-level and most useful filter of all.

ARCHITECTURE

5 Netpage System Architecture

5.1 THE INTERNET

The Internet is a worldwide collection of interconnected networks which communicate using the TCP/IP protocol suite [89]. A TCP/IP-based internetwork not connected to *the* Internet is often referred to as *an* internet (i.e. with a lower-case 'i'). When an internet is deployed within an organization, it is often termed an intranet.

Access to the Internet is widespread in developed countries. In the United States, for example, 41% of the population has access to the Internet [78].

While most consumers still access the Internet via low-speed dial-up modems connected to the switched telephone system, inexpensive broadband access is becoming available to a growing number of households via the cable networks. Cable service is available to 95% of American households [87], and cable modem service is available to a 20% subset [8]. While dial-up modems offer speeds of up to 56Kbps, cable modems offer practical speeds of up to about 3Mbps¹, i.e. over 50 times faster.

DSL (Digital Subscriber Line) [11,12], while offering similar speeds to cable modems but via the telephone system, is not yet widely used. ISDN (Integrated Services Digital Network), although widely used for corporate access, has had little consumer impact due to its high price and comparatively low performance.

The deployment of third-generation (3G) cellular telephony within the next few years will bring practical mobile broadband speeds of 2Mbps [25]. 3G cellular uses WCDMA (wide-band code-division multiple access), a spread-spectrum technology. Satellite systems are arguably closer to offering even faster broadband Internet access [39,72,93].

The core of the Internet is made up of a number of independent high-speed fiber-optic networks connected into NAPs (Network Access Points) or peered directly. These have until recently used single-wavelength TDM (Time-Division Multiplexing) SONET (Synchronous Optical Network) transmission systems which utilize about 1% of an optic fiber's capacity to yield a 2.5Gbps OC-48 channel². Carriers are now beginning to deploy multi-wavelength DWDM (Dense Wavelength-Division Multiplexing) systems which yield up to 40 such channels per optic fiber, thus increasing network capacity significantly without requiring the laying of more fiber [9,53]. Internet architects are therefore now contemplating aggregate capacity in the terabit (Tbps) range.

The Internet uses the four-layer TCP/IP protocol suite. The application layer provides various end-to-end application services, and is a client of the transport layer which provides end-to-end delivery services. The transport layer in turn is a client of the network layer which provides packet routing. The network layer is a client of the link layer which encapsulates specifics of the protocols and hardware of the actual communications links.

The core Internet transport protocol, TCP (Transmission Control Protocol), provides a reliable end-to-end delivery service. The core Internet network protocol, IP (Internet Protocol), provides an unreliable and connectionless packet routing service. IP may lose or

1. Although the cable supports 30Mbps and the cable modem theoretically supports 10Mbps.

2. SONET channels have an OC-*n* designation, where OC stands for Optical Carrier, and *n* gives the channel speed in units of about 52Mbps. An OC-48 channel therefore has a speed of about 2.5Gbps.

deliberately discard packets, and may deliver packets out of order, and it is the responsibility of a higher layer to provide a reliable end-to-end service.

With the proliferation of streaming media services on the Internet, support for multicast is spreading rapidly. Multicast is a form of broadcast with a specific set of recipients. It makes efficient use of network capacity because a packet traverses a network link once rather than once per recipient. It is particularly efficient if the recipients are connected to the Internet via an intrinsically broadcast medium such as cable or satellite. The @Home cable network has successfully enabled multicasting of streaming media services [50].

IP Multicast is an extension of IP, and so is unreliable. While this is often acceptable for time-critical data such as streaming video, it may not be acceptable for other shared data types. Significant effort is being expended to develop reliable multicast transport protocols on top of IP Multicast. Although several reliable multicast protocols are available and have been deployed [56,57,40,41], the Internet standardization process is incomplete [42].

5.2 NETPAGES AND NETPAGE DOCUMENTS

Netpages are the foundation on which a Netpage Network is built. They provide a paper-based user interface to published information and interactive services.

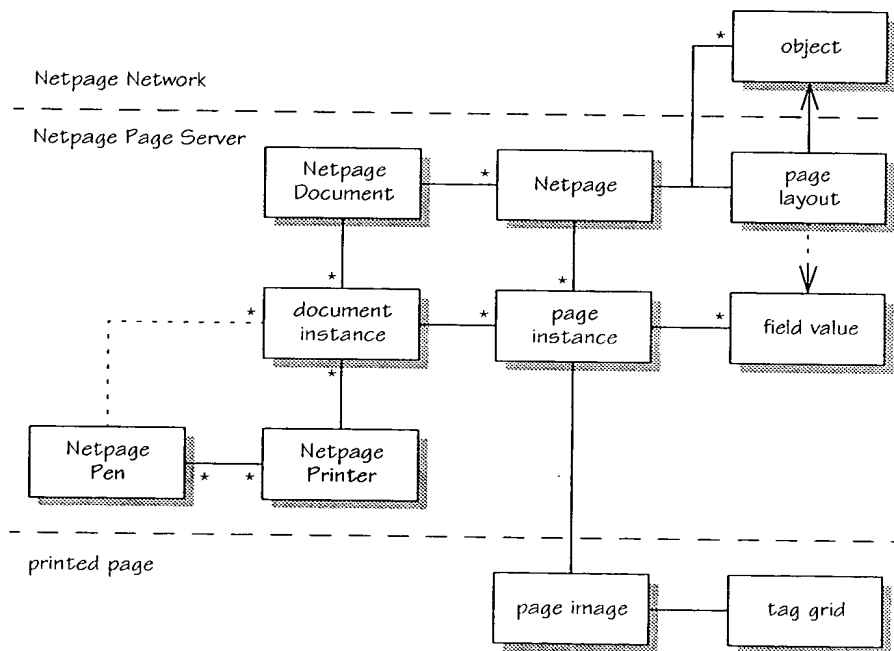


Figure 1. Netpage Document structure

Each Netpage consists of a compact page layout maintained persistently by a Netpage Page Server. The page layout refers to objects such as images, fonts and pieces of text, typically stored elsewhere on the Netpage Network.

Netpages are organized into Netpage Documents. Both Netpages and Netpage Documents are assigned globally unique identifiers.

Each Netpage Document has a set of document instances, each of which describes a printed instance of the document. Each Netpage in the Netpage Document has a corresponding set of page instances, each of which describes a printed instance of the page. Both page instances and document instances are assigned globally unique identifiers. They are also uniquely associated with the printer on which they are printed, and the pen which initiated the print request, if known.

Each page instance maintains a set of user-supplied values for fields in the page layout. This ensures that user input is captured and stored independently for each page instance. The separation of page instances and Netpages is crucial for pages which contain input fields, i.e. forms. It is not crucial for pages devoid of input fields, but still useful because it supports independent auditing of each page instance.

The physical page image includes encoded information which identifies the page instance and hence the Netpage to which it corresponds. It also includes encoded information which superimposes an addressable spatial grid over the page image, to allow pen actions performed relative to the page image to be correlated with the contents of the page layout.

The encoded information is normally printed in infrared-absorptive ink on any normal paper substrate which is infrared-reflective. Near-infrared wavelengths are invisible to the human eye but are easily sensed by a solid-state image sensor with an appropriate filter.

The encoded information is picked up by an infrared-imaging pen and transmitted to the associated Netpage Printer. The pen is wireless and communicates with the Netpage Printer using an encrypted radio frequency signal.

The encoded information is organized as a set of tags, each containing both the id of the page instance and the position of the tag. The tags tile the entire page image, and are sufficiently small and densely arranged that the pen can reliably image at least one tag even on a single click on the page. It is important that the pen recognize the page instance id and position on every interaction with the page, since the interaction is stateless.

The tags are error-correctably encoded to make them resilient to errors introduced by dirt on the page or during the imaging process.

Memjet-based Netpage Printers are designed to print a tag grid using infrared (IR) ink. Printers not enabled for IR printing have the option to print tags using IR-absorptive black ink, although this restricts tags to otherwise empty areas of the page. Although such pages have more limited functionality than IR-printed pages, they are still classed as Netpages.

5.3 THE NETPAGE NETWORK

A Netpage Network consists of a distributed set of Netpage Publication Servers, Netpage Page Servers, and Netpage Printers connected via an internet. In technological terms this document describes *any* Netpage Network. In business terms it usually refers to *the* Netpage Network connected via *the* Internet.

As indicated above, a Netpage Page Server maintains persistent information about Netpage Documents, Netpages, and their printed instances, to allow pen operations on printed pages to be interpreted intelligently.

The Netpage Network includes any number of Netpage Page Servers, each handling a subset of Netpages. As indicated above, each page instance is identified by a globally unique

id which is encoded in the tag grid of the corresponding printed page. The Netpage Printer uses this id to retrieve the page layout of the page from a Netpage Page Server when it needs to interpret pen operations relative to the page.

The Netpage Printer uses the internet Distributed Name System (DNS) to resolve a Netpage instance id into a page instance maintained by a particular Netpage Page Server.

The DNS is a protocol and a hierarchical system of name servers used to resolve internet domain names into resources. Planned enhancements to the DNS allow it to be used to resolve more general Uniform Resource Identifiers (URIs), and in particular Uniform Resource Names (URNs), into resource locations [22]. Netpage instance ids are formulated as URNs, allowing the enhanced DNS to be used to resolve them. In the absence of timely standardization and deployment of an enhanced DNS on the Internet, the Netpage Network can deploy its own system of enhanced name servers.

A Netpage Publication Server is an internet server which publishes Netpage Documents to Netpage Printers. It is described in later sections.

5.4 THE NETPAGE PRINTER

The Netpage Printer is the appliance which prints Netpage Documents. It is connected to a Netpage Network via an internet, ideally via a broadband connection.

Apart from identity and security settings in non-volatile memory, the Netpage Printer contains no persistent storage. As far as a user is concerned, *the network is the computer* [91]. Netpages function interactively across space and time with the help of the distributed Netpage Page Servers, independently of particular Netpage Printers.

The Netpage Printer receives Netpage Documents from Netpage Publication Servers. Each document is distributed in two parts: the page layouts, and the actual text and image objects which populate the pages. Because of personalization, page layouts are typically specific to a particular subscriber and so are pointcast to the subscriber's printer. Text and image objects, on the other hand, are typically shared with other subscribers, and so are multicast to all subscribers' printers.

The Netpage Publication Server optimizes the segmentation of document content into pointcasts and multicasts. After receiving the pointcast of a document's page layouts, the printer knows which multicasts, if any, to listen to.

Once the printer has received the entire document's page descriptions, i.e. page layouts *and* objects, it can print the document.

The printer rasterizes and prints odd and even pages simultaneously on both sides of the sheet. It therefore contains duplexed print engines and imaging units.

The printing process consists of two decoupled stages: rasterization of page descriptions, and expansion and printing of page images. The raster image processor (RIP) consists of one or more standard DSPs running in parallel. The duplexed print engines consist of custom processors which expand, dither and print page images in real time, synchronized with the operation of the printheads in the imaging units.

There are four major design variations embodied in the various Netpage Printer models:

- *form factor*: portable, wall-mount or tabletop
- *printhead width*: 8½" (portrait Letter) or 11" (landscape Letter)
- *paper source*: cut sheet or paper roll
- *Internet connection*: wired or wireless

This form factor variations yield three basic models, each with variants determined by printhead width (and hence printing speed), and paper source. Six planned models are defined in Table 1, and illustrated in Figure 2.

Table 1. Netpage Printer models

model	form factor	variant	printhead width	paper source	Internet connect
Netpage Travelprinter	portable	R	8½"	paper roll	wireless
Netpage Wallprinter	wall-mount	-	11"	cut sheet	wired or wireless
		Pro		cut sheet	
		Pro R		paper roll	
Netpage Tableprinter	tabletop	Pro	11"	cut sheet	
		Pro R		paper roll	




Figure 2. Netpage Printer family

The Wallprinter and Tableprinter models can be factory-configured with various network modules, allowing both wired and wireless versions. The Travelprinter uses a cellular telephone module, with the promise of broadband speed within a few years.

The Wallprinter models are ideal for unobtrusive installation in a home, while the Tableprinter models might be preferred in an office environment. Note that the Tableprinter models are Wallprinter models factory-adapted for tabletop use via a stand.

The paper roll cartridge contains both paper and ink. The paper is in the form of a continuous roll, cut on demand by the printer. The 11" paper cartridge has a capacity of 1000 Letter sheets. It also contains the glue supply for binding the sheets of a document together. The 8½" paper cartridge has a capacity of 50 Letter sheets, or equivalently 100 A5 sheets. It doesn't contain a glue supply because the Travelprinter doesn't include a binding mechanism.

The 8½" printhead models print at 60 pages per minute (i.e. 30 sheets per minute). The 11" printhead models print at 90 pages per minute (i.e. 45 sheets per minute).

5.5 THE NETPAGE PEN

The Netpage Pen operates both as a normal marking ink pen and as a non-marking stylus. When either nib is in contact with a Netpage, the pen continuously monitors its movements relative to the page. The nib is attached to a pressure sensor. The pen pressure can be interpreted relative to a threshold to indicate whether the pen is "up" or "down". It can also be interpreted as a continuous value, for example when the pen is capturing a signature, to allow the full dynamics of the signature to be verified.

The pen determines the position of its nib on the Netpage by imaging, in the infrared spectrum, an area of the page in the vicinity of the nib. It decodes the nearest page id and position tag, and adjusts the position given by the tag to account for the distance between the area imaged and the actual nib, and the position of the tag in the imaged area. Although the position resolution of the tag may be low, because the tag density on the page is inversely proportional to the tag size, the adjusted position resolution is quite high, and easily exceeds the minimum 200 dpi resolution required for handwriting recognition [92].

Pen actions relative to a Netpage consist of a series of strokes. A stroke consists of a sequence of time-stamped pen positions on the page, initiated by a pen-down event and completed by the subsequent pen-up event. A stroke is also tagged with the page id of the Netpage whenever the page id changes, i.e. just at the start of the stroke under normal circumstances.

The position tags on the Netpage contain various control bits. One of these instructs the pen to activate its "active area" LED. Thus a region on the page which corresponds to the active area of a button or hyperlink can be encoded to activate this LED, giving the user visual feedback that the button or hyperlink is active when the pen passes over it. Another control bit instructs the pen to capture continuous pen pressure readings and tag the stroke with these readings. Thus a region on the page which corresponds to a signature input area can be encoded to capture continuous pen pressure.

Whenever the pen is within range of a printer with which it can communicate, the pen slowly flashes its "online" LED. When the pen fails to decode a stroke relative to the page, it momentarily activates its "error" LED. When the pen succeeds in decoding a stroke relative to the page, it momentarily activates its "ok" LED.

The pen also contains a pair of passive accelerometers mounted at right angles to each other in the plane normal to the pen's axis. The accelerometers respond to gravity and allow the pen to compute its tilt. This in turn helps it auto-focus its optics and compute the nib-to-tag displacement. If the stroke is being tagged with pen pressure readings, then it is also tagged with tilt readings.

A sequence of captured strokes, whether tagged with pen pressure and tilt or not, is referred to as *digital ink*. Digital ink forms the basis for the digital exchange of drawings and handwriting, for on-line recognition of handwriting [92], and for on-line verification of signatures.

The pen is wireless and transmits digital ink to the Netpage Printer using a radio frequency signal. The digital ink data is encrypted for security and packetized for efficient transmission, but is always flushed on a pen-up event to ensure timely handling in the printer.

When the pen is out-of-range of a printer it buffers digital ink in internal memory, which has a capacity of more than 12 minutes of continuous handwriting. When the pen is once again within range of a printer, it transfers any buffered digital ink.

A pen can be registered with any number of printers, but because all state data resides in Netpages both on paper and on the network, it is immaterial which printer a pen is communicating with at any particular time.

5.6 NETPAGE INTERACTION

When the Netpage Printer receives a digital ink stroke from the pen, it retrieves the page layout of the Netpage identified in the stroke, to allow it to correctly interpret the stroke. The printer resolves, via the DNS, the address of the Netpage Page Server which holds the page layout, and then retrieves the page layout from the server. If the page was recently identified in an earlier stroke, then the printer may already have the address of the relevant Netpage Page Server in its cache. It may also have the page layout itself in its cache, in which case there may be no need to retrieve it.

Once the printer has the page layout of the Netpage to which the pen stroke refers, it can interpret the stroke in relation to the layout and content of the page. This involves hit-testing the objects on the page to determine which objects the pen is interacting with, in much the same way that mouse movements and button presses are interpreted in a graphical user interface system.

A "click" is a stroke where the distance between the pen down position and the subsequent pen up position is less than some small maximum. An object which is activated by a click requires a click to be activated, i.e. a longer stroke is ignored. The failure of a pen action, such as a "sloppy" click, to register is indicated by the lack of response from the pen's "ok" LED.

There are two kinds of interactive objects on a Netpage: hyperlinks and form fields.

When a hyperlink is activated, the printer sends a request to a handler somewhere on the network. The handler is identified by a URI, and the URI is resolved in the normal way via the DNS. There are three types of hyperlinks: general hyperlinks, form hyperlinks, and selection hyperlinks. A general hyperlink may implement a request for a linked document, or may simply signal a preference to a server. A form hyperlink submits the corresponding form to a form handler. A selection hyperlink submits the current selection to a selection

handler. If the current selection contains a single-word piece of text, for example, the selection handler may return a single-page document giving the word's meaning within the context in which it appears, or a translation into a different language. Each hyperlink type is characterized by what information is submitted to the handler.

Form fields come in four varieties: checkboxes, text areas, digital ink areas, and signature areas. A checkbox accepts a true or false value. Any tick or short stroke captured in a checkbox area is assigned as a true value to the corresponding field. A text area accepts a text string. Any digital ink captured in a text area is automatically converted to text via on-line handwriting recognition and the text is assigned to the corresponding field. A digital ink area accepts raw digital ink. Any digital ink captured in a digital ink area is assigned to the corresponding field. A signature area accepts a handwritten signature. Any digital ink captured in a signature area is automatically verified and the resulting signature token is assigned to the corresponding signature field. Signature verification is discussed in detail in a later section.

Because the handwriting recognition algorithm works "on-line" (i.e. with access to the dynamics of the pen movement), rather than "off-line" (i.e. with access only to a bitmap of pen markings), it can recognize run-on discretely-written characters [92] with high accuracy, without a writer-dependent training phase.

Digital ink, as already stated, consists of a sequence of strokes. Any stroke which starts in a particular object's active area is appended to that area's digital ink stream, ready for eventual interpretation. Any stroke not appended to the remaining inactive area's digital ink stream.

Digital ink captured in the inactive area is interpreted as a selection gesture. Any circumscription of an object is interpreted as a selection of that object.

The printer maintains a current selection for each pen. The selection contains the most recent object selected, resolved with reference to the page layout and content. The selection can be attached to or pasted into another form, or in general be submitted to a selection handler as described earlier. The selection is cleared after an inactivity time-out.

Table 2 provides a summary of pen interactions with a Netpage.

Table 2. Summary of pen interactions with a Netpage

object	type	pen input	action
hyperlink	general	click	submit action to handler via URI
	form	click	submit form to handler via URI
	selection	click	submit selection to handler via URI
form field	checkbox	tick ^a	set field value to true
	text area	handwriting	convert digital ink to text; assign text to field
	digital ink area	digital ink	assign digital ink to field
	signature area	signature	verify digital ink signature; assign signature token to field
none	-	circumscription	convert digital ink to region; select object(s) in region

a. i.e. any short stroke

As described earlier, user input on a physical Netpage is ultimately recorded persistently by a Netpage Page Server together with the corresponding page instance. To ensure efficient capture of user input, the printer accumulates input locally. To prevent update anomalies, however, the printer temporarily obtains exclusive access to the page instance from the Netpage Page Server. The printer flushes input back to the server and relinquishes exclusive access when the user initiates a non-local action on the page; after an inactivity time-out on the page; when the printer wishes to free up local storage consumed by the page; and on request from the server.

When the printer submits a form to a form handler, it simply submits the document instance of the form. The form handler retrieves the field values from the Netpage Page Server at its leisure.

A form can also act as a shared “blackboard” between the user and the form handler, i.e. the form handler can query the contents of the form fields maintained by the Netpage Page Server without the user explicitly submitting the form.

For text areas, the raw digital ink is optionally also stored with the page instance on the Netpage Page Server. This allows the form handler to interrogate the raw digital ink should it suspect the original recognition of the handwriting. This might involve human intervention at the application level for forms which fail certain application-specific consistency checks.

Form fields can optionally be tagged to indicate their meaning. Fields tagged in this way may include name and address fields, for example. This semantic tagging allows these fields to be automatically filled in whenever a “blank” form is requested by an identifiable user, i.e. a user who has registered their identity with the system and linked it to the identity of their pen.

5.7 STANDARD FEATURES OF NETPAGES

Each Netpage is printed with the Netpage logo at the bottom to indicate that it is a Netpage and therefore has interactive properties. The logo also acts as a “copy” button. In most cases pressing the logo produces a copy of the page. In the case of a form the button instead elicits a page giving the user the option to print the entire form document. And in the case of a secure document, such as a ticket or coupon, the button elicits an explanatory note or advertising page.

The default single-page copy function is handled directly by the relevant Netpage Page Server. Special copy functions are handled by linking the logo button to other URIs.

Once a Netpage form has been submitted, it is marked as submitted by the Netpage Page Server and cannot be submitted again. An attempt to do so elicits a status report indicating when it was submitted. A copy of the form can still be made, altered, and re-submitted.

5.8 THE HELP SYSTEM

The Netpage Printer has a single button labelled “help”. When pressed it elicits a single page of information. This information includes the following:

- status of printer connection
- status of printer consumables
- top-level help menu

- document function menu
- top-level Netpage Network directory

The help menu provides a hierarchical manual on how to use the Netpage System.

The document function menu includes the following functions:

- print a copy of a document
- print a clean copy of a form
- print the status of a document

A document function is initiated by simply pressing the button and then touching any page of the document. The status of a document indicates who published it and when, to whom it was delivered, and to whom and when it was subsequently submitted as a form.

The Netpage Network directory allows the user to navigate the hierarchy of publications and services on the network.

The help page is obviously unavailable if the printer is unable to print. In this case the "error" light is lit and the user can request remote diagnosis over the network.

6 Personalized Publication Model

In the following discussion, news is used as a canonical publication example to illustrate personalization mechanisms in the Netpage System. Although news is often used in the limited sense of newspaper and news magazine news, the intended scope is wider.

In the Netpage System, the editorial content and the advertising content of a news publication are personalized using different mechanisms. The editorial content is personalized according to the reader's explicitly stated and implicitly captured interest profile. The advertising content is personalized according to the reader's locality and demographic.

6.1 EDITORIAL PERSONALIZATION

A subscriber can draw on two kinds of news sources: those that deliver news publications, and those that deliver news streams. While news publications are aggregated and edited by the publisher, news streams are aggregated either by a news publisher or by a specialized news aggregator. News publications typically correspond to traditional newspapers and news magazines, while news streams can be many and varied: a "raw" news feed from a news service, a cartoon strip, a freelance writer's column, a friend's bulletin board, or the reader's own e-mail.

The Netpage Publication Server supports the publication of edited news publications as well as the aggregation of multiple news streams. By handling the aggregation and hence the formatting of news streams selected directly by the reader, the server is able to place advertising on pages over which it otherwise has no editorial control.

The subscriber builds a daily newspaper by selecting one or more contributing news publications, and creating a personalized version of each. The resulting daily editions are printed and bound together into a single newspaper. The various members of a household typically express their different interests and tastes by selecting different daily publications and then customizing them.

For each publication, the reader optionally selects specific sections. Some sections appear daily, while others appear weekly. The daily sections available from The New York Times online, for example, include "Page One Plus", "National", "International", "Opinion", "Business", "Arts/Living", "Technology", and "Sports". The set of available sections is obviously specific to a publication, as is the default subset.

The reader extends the daily newspaper by creating custom sections, each one drawing on any number of news streams. Custom sections might be created for e-mail and friends' announcements ("Personal"), or for monitoring news feeds for specific topics ("Alerts" or "Clippings").

For each section, the reader optionally specifies its size, either qualitatively (e.g. short, medium, or long), or numerically (i.e. as a limit on its number of pages), and the desired proportion of advertising, either qualitatively (e.g. high, normal, low, none), or numerically (i.e. as a percentage).

The reader also optionally expresses a preference for a large number of shorter articles or a small number of longer articles. Each article is ideally written (or edited) in both short and long forms to support this preference.

An article may also be written (or edited) in different versions to match the expected sophistication of the reader, for example to provide children's and adults' versions. The appropriate version is selected according to the reader's age.

The articles which make up each section are selected and prioritized by the editors, and each is assigned a useful lifetime. By default they are delivered to all relevant subscribers, in priority order, subject to space constraints in the subscribers' editions.

In sections where it is appropriate, the reader may optionally enable collaborative filtering. This is then applied to articles which have a sufficiently long lifetime. Each article which qualifies for collaborative filtering is printed with rating buttons at the end of the article. The buttons can provide an easy choice (e.g. "liked" and "disliked"), making it more likely that readers will bother to rate the article.

Articles with high priorities and short lifetimes are therefore effectively considered essential reading by the editors and are delivered to most relevant subscribers.

The reader optionally specifies a serendipity factor, either qualitatively (e.g. do or don't surprise me), or numerically. A high serendipity factor lowers the threshold used for matching during collaborative filtering. A high factor makes it more likely that the corresponding section will be filled to the reader's specified capacity. A different serendipity factor can be specified for different days of the week.

The reader also optionally specifies topics of particular interest within a section, and this modifies the priorities assigned by the editors.

The speed of the reader's Internet connection affects the quality at which images can be delivered. The reader optionally specifies a preference for fewer images or smaller images or both. If the number or size of images is not reduced, then images may be delivered at lower quality.

At a global level, the reader specifies how quantities, dates, times and monetary values are localized. This involves specifying whether units are imperial or metric, a local timezone and time format, and a local currency, and whether the localization consist of *in situ* translation or annotation.

To reduce reading difficulties caused by poor eyesight, the reader optionally specifies a global preference for a larger presentation. Both text and images are scaled accordingly, and less information is accommodated on each page.

The language in which a news publication is published, and its corresponding text encoding, is a property of the publication and not a preference expressed by the user. However, the Netpage Network may provide automatic translation services in various guises.

6.2 ADVERTISING LOCALIZATION AND TARGETING

The personalization of the editorial content directly affects the advertising content, because advertising is typically placed to exploit the editorial context. Travel ads, for example, are more likely to appear in a travel section than elsewhere. The value of the editorial content to an advertiser (and therefore to the publisher) lies in its ability to attract large numbers of readers with the right demographics.

Effective advertising is placed on the basis of locality and demographics. Locality determines proximity to particular services, retailers etc., and particular interests and concerns associated with the local community and environment. Demographics determine general interests and preoccupations as well as likely spending patterns.

A news publisher's most profitable product is advertising "space", a multi-dimensional entity determined by the publication's geographic coverage, the size of its readership, its readership demographics, and the page area available for advertising.

In the Netpage System, the Netpage Publication Server computes the approximate multi-dimensional size of a publication's saleable advertising space on a per-section basis, taking into account the publication's geographic coverage, the section's readership, the size of each reader's section edition, each reader's advertising proportion, and each reader's demographic.

In comparison with other media, the Netpage System allows the advertising space to be defined in greater detail, and allows smaller pieces of it to be sold separately. It therefore allows it to be sold at closer to its true value.

For example, the same advertising "slot" can be sold in varying proportions to several advertisers, with individual readers' pages randomly receiving the advertisement of one advertiser or another, overall preserving the proportion of space sold to each advertiser.

The Netpage System allows advertising to be linked directly to detailed product information and online purchasing. It therefore raises the intrinsic value of the advertising space.

Because personalization and localization are handled automatically by Netpage Publication Servers, an advertising aggregator can provide arbitrarily broad coverage of both geography and demographics. The subsequent disaggregation is efficient because it is automatic. This makes it more cost-effective for publishers to deal with advertising aggregators than to directly capture advertising. Even though the advertising aggregator is taking a proportion of advertising revenue, publishers may find the change profit-neutral because of the greater efficiency of aggregation. The advertising aggregator acts as an intermediary between advertisers and publishers, and may place the same advertisement in multiple publications.

It is worth noting that ad placement in a Netpage publication can be more complex than ad placement in the publication's traditional counterpart, because the publication's advertising space is more complex. While ignoring the full complexities of negotiations between advertisers, advertising aggregators and publishers, it is clear that the Netpage System should ideally provide some automated support for these negotiations, including support for automated auctions of advertising space. Automation is particularly desirable for the placement of advertisements which generate small amounts of income, i.e. small or highly localized advertisements.

Once placement has been negotiated, the aggregator captures and edits the advertisement and records it on a Netpage Ad Server. Correspondingly, the publisher records the ad placement on the relevant Netpage Publication Server. When the Netpage Publication Server lays out each user's personalized publication, it picks the relevant advertisements from the Netpage Ad Server.

6.3 USER PROFILES

The personalization of news and other publications relies on an assortment of user-specific profile information:

- publication customizations
- collaborative filtering vectors
- contact details
- presentation preferences

The customization of a publication is typically publication-specific, and so the customization information is maintained by the relevant Netpage Publication Server.

A collaborative filtering vector consists of the user's ratings of a number of news items. As described earlier, it is used to correlate different users' interests for the purposes of making recommendations. Although there are benefits to maintaining a single collaborative filtering vector independently of any particular publication, there are two reasons why it is more practical to maintain a separate vector for each publication: there is likely to be more overlap between the vectors of subscribers to the same publication than to different publications; and a publication is likely to want to present its users' collaborative filtering vectors as part of the value of its brand, not to be found elsewhere. Collaborative filtering vectors are therefore also maintained by the relevant Netpage Publication Server.

Contact details, including name, street address, ZIP code, state, country, telephone numbers, etc., are by their nature global and are maintained by a Netpage Registration Server.

Presentation preferences, including those for quantities, dates and times discussed earlier, are likewise global and maintained in the same way.

The localization of advertising relies on the locality indicated in the user's contact details, while the targeting of advertising relies on personal information such as date of birth, gender, marital status, income, profession, education, etc., or qualitative derivatives such as age range and income range.

For those users who choose to reveal personal information for advertising purposes, the information is maintained by the relevant Netpage Registration Server. In the absence of such information, advertising can be targeted on the basis of the demographic associated with the user's ZIP or ZIP+4 code.

Each user, pen and printer is assigned a globally unique identifier, and the Netpage Registration Server maintains the relationships between them. The server also keeps track of which publications a user has authorized to print on particular printers. Each user may have several pens, but a pen is specific to a single user. A pen may know any number of

printers, and a printer may know any number of pens. These relationships are illustrated in Figure 3.




Figure 3. User registration relationships

The pen identifier is used, in the form of a URN, to locate the corresponding user profile maintained by a particular Netpage Registration Server, via the DNS in the usual way.

6.4 INTELLIGENT PAGE LAYOUT

The Netpage Publication Server automatically lays out the pages of each user's personalized publication on a section-by-section basis. Since most advertisements are in the form of pre-formatted rectangles, they are placed on the page before the editorial content.

The advertising ratio for a section can be achieved with wildly varying advertising ratios on individual pages within the section, and the ad layout algorithm exploits this. The algorithm attempts to co-locate closely tied editorial and advertising content, e.g. ads for roofing material placed specifically with the publication because of a special feature on do-it-yourself roofing repairs.

The editorial content selected for the user, i.e. text with associated images and graphics, is then laid out according to various aesthetic rules.

The entire process, including the selection of ads and the selection of editorial content, must be iterated once the layout has converged, to attempt to more closely achieve the user's stated section size preference. The section size preference can, however, be matched *on average* over time, allowing significant day-to-day variations.

6.5 DOCUMENT FORMAT

Once the document is laid out, it is encoded for efficient distribution and persistent storage on the Netpage Network.

The primary efficiency mechanism is the separation of information specific to a single user's edition and information shared between multiple users' editions. The specific information consists of the page layout. The shared information consists of the objects to which the page layout refers, including images, graphics, and pieces of text.

A text object contains fully-formatted text represented in the Extensible Markup Language (XML) [101] using the Extensible Stylesheet Language (XSL) [102]. XSL provides

precise control over text formatting independently of the region into which the text is being set, which in this case is being provided by the layout. The text object contains embedded language codes to enable automatic translation, and embedded hyphenation hints to aid with paragraph formatting.

An image object encodes an image in the JPEG 2000 [44,45] wavelet-based compressed image format to achieve high quality at compression ratios exceeding 10:1.

A graphic object encodes a 2D graphic in Scalable Vector Graphics (SVG) [104] format.

The layout itself consists of a series of placed image and graphic objects, linked textflow objects through which text objects flow, hyperlinks and input fields as described earlier, and watermark regions. These layout objects are summarized in Table 3. The layout uses a compact format suitable for efficient distribution and storage.

The layout is tagged with the version of the text-setting algorithm used by the Netpage Publication Server when the layout was first created, allowing the Netpage Printer to exactly reproduce the physical layout intended by the server.

Because Netpage Printer software is automatically upgraded over the Netpage Network, it is feasible to for Netpage Printers to contain every version of the text-setting algorithm.

Table 3. Netpage layout objects

layout object	attribute	linked object format
image	position	-
	image object URI	JPEG 2000
graphic	position	-
	graphic object URI	SVG
textflow	textflow id	-
	region ^a	
	optional text object URI	XML/XSL
hyperlink	type	-
	region ^a	-
	handler URI	-
field	type	-
	meaning	
	region ^a	-
watermark	region ^a	-

a. arbitrary multi-edged shape defined with spline paths

6.6 DOCUMENT DISTRIBUTION

As described above, for purposes of efficient distribution and persistent storage on the Netpage Network, a user-specific page layout is separated from the shared objects to which it refers.

When a subscribed publication is ready to be distributed, the Netpage Publication Server allocates, with the help of the Netpage Id Server, a globally unique id for each page, page instance, document, and document instance.

The server computes a set of optimized subsets of the shared content and creates a multicast channel for each subset, and then tags each user-specific layout with the names of the multicast channels which will carry the shared content used by that layout. The server then pointcasts each user's layouts to that user, and when the pointcasting is complete, multicasts the shared content on the specified channels. After receiving its pointcast, each Netpage Printer subscribes to the multicast channels specified in the page layouts. During the multicasts, each printer extracts from the multicast streams those objects referenced by its page layouts.

Once a printer has received all the objects to which its page layouts refer, the printer re-creates the fully-populated layout and then rasterizes and prints it.

The server also delivers each page layout to the relevant Netpage Page Server, which may be co-located with the Netpage Publication Server, or may be located elsewhere on the network. Thus the page layouts are persistently archived as Netpages. It is the responsibility of the Netpage Publication Server to preserve the shared objects referenced by the Netpages, to ensure that they are really persistent. It may choose to archive these shared objects elsewhere on the network at any time. The object URIs embedded in the Netpages allow the objects to move.

Because of limited memory in a Netpage Printer, the printer may be unable to make use of a multicast when it first occurs. The Netpage Publication Server therefore allows printers to submit requests for re-multicasts. When a critical number of requests is received, the server re-multicasts the corresponding shared objects.

Once a document is printed, a Netpage Printer can produce an exact duplicate at any time by retrieving its page layouts from the relevant Netpage Page Server, and retrieving the objects to which it refers from the network.

7 Security

7.1 CRYPTOGRAPHY

Cryptography is used to protect sensitive information, both in storage and in transit, and to authenticate parties to a transaction. There are two classes of cryptography in widespread use: secret-key cryptography and public-key cryptography. The Netpage Network uses both classes of cryptography.

Secret-key cryptography, also referred to as symmetric cryptography, uses the same key to encrypt and decrypt a message. Two parties wishing to exchange messages must first arrange to securely exchange the secret key.

Public-key cryptography, also referred to as asymmetric cryptography, uses two encryption keys. The two keys are mathematically related in such a way that any message encrypted using one key can only be decrypted using the other key. One of these keys is then published, while the other is kept private. The public key is used to encrypt any message intended for the holder of the private key. Once encrypted using the public key, a message can only be decrypted using the private key. Thus two parties can securely exchange messages without first having to exchange a secret key. To ensure that the private key is secure, it is normal for the holder of the private key to generate the key pair.

Public-key cryptography can be used to create a digital signature. If the holder of the private key creates a known hash of a message and then encrypts the hash using the private key, then anyone can verify that the encrypted hash constitutes the "signature" of the holder of the private key with respect to that particular message, simply by decrypting the encrypted hash using the public key and verifying the hash against the message. If the signature is appended to the message, then the recipient of the message can verify both that the message is genuine and that it has not been altered in transit.

To make public-key cryptography work, there has to be a way to distribute public keys which prevents impersonation. This is normally done using certificates and certificate authorities. A certificate authority is a trusted third party which authenticates the connection between a public key and someone's identity. The certificate authority verifies the person's identity by examining identity documents etc., and then creates and signs a digital certificate containing the person's identity details and public key. Anyone who trusts the certificate authority can use the public key in the certificate with a high degree of certainty that it is genuine. They just have to verify that the certificate has indeed been signed by the certificate authority, whose public key is well-known.

In most transaction environments, public-key cryptography is only used to create digital signatures and to securely exchange secret session keys. Secret-key cryptography is used for all other purposes.

In the following discussion, when reference is made to the *secure* transmission of information between a Netpage Printer and a server, what actually happens is that the printer obtains the server's certificate, authenticates it with reference to the certificate authority, uses the public key-exchange key in the certificate to exchange a secret session key with the server, and then uses the secret session key to encrypt the message data. A *session* key, by definition, can have an arbitrarily short lifetime.

7.2 NETPAGE PRINTER SECURITY

Each Netpage Printer is assigned a pair of unique identifiers at time of manufacture which are stored in read-only memory in the printer and in the Netpage Registration Server database. The first id is public and uniquely identifies the pen on the Netpage Network. The second id is secret and is used when the printer is first registered on the network.

When the printer connects to the Netpage Network for the first time after installation, it creates a signature public/private key pair. It transmits the secret id and the public key securely to the Netpage Registration Server. The server compares the secret id against the printer's secret id recorded in its database, and accepts the registration if the ids match. It then creates and signs a certificate containing the printer's public id and public signature key, and stores the certificate in the registration database.

The Netpage Registration Server acts as a certificate authority for Netpage Printers, since it has access to secret information allowing it to verify printer identity.

When a user subscribes to a publication, a record is created in the Netpage Registration Server database authorizing the publisher to print on a particular printer. Every document sent to a printer is signed by the publisher using the publisher's private signature key. The printer verifies via the registration database that the publisher is authorized to print on the printer, and verifies the signature using the publication's public key, obtained from the publisher's certificate stored in the registration database.

The Netpage Registration Server accepts requests to add printing authorizations to the database, so long as those requests are initiated via a pen registered to the printer.

The user can register a Web terminal as a "publisher" authorized to print on a printer. This is useful if the user has a Web terminal in the home which is used to locate documents on the Web for printing. The authorization occurs as follows: the user prints a Web terminal authorization form. The Netpage Registration Server generates a short-lifetime one-time-use id for the Web terminal which is printed on the form, together with the URI of the printer. The Web terminal is used to navigate to a Netpage Registration Server registration site, where the one-time-use id is entered, as well as the URI of the printer. The Web terminal generates a signature public/private key pair. The server allocates a publisher id for the Web terminal, creates and signs a certificate containing the publisher id and the public key, and stores the certificate in the registration database. The URI of the printer, the publisher id, and the private signature key are stored locally on the Web terminal.

Whenever the Web terminal wishes to print on the printer, it sends the printer a request containing the URI of the document to be printed, together with the publisher id, signed with the Web terminal's private signature key. On receipt of the request and before acting on it, the printer verifies publisher id and signature in the usual way.

The user can print a list of current printing authorizations at any time, and retract any which are being abused.

7.3 NETPAGE PEN SECURITY

Each Netpage Pen is assigned a unique identifier at time of manufacture which is stored in read-only memory in the pen and in the Netpage Registration Server database. The id uniquely identifies the pen on the Netpage Network.

A Netpage Pen can know a number of Netpage Printers, and a printer can know a number of pens. A pen communicates with a printer via a radio frequency signal whenever it is within range of the printer. Once a pen and printer are registered, they regularly exchange session keys. Whenever the pen transmits digital ink to the printer, it always uses the appropriate session key. Digital ink is never transmitted in the clear.

A pen stores a session key for every printer it knows, indexed by printer id, and a printer stores a session key for every pen it knows, indexed by pen id. Both have a finite storage capacity for session keys, and will forget a session key on a least-recently-used basis if necessary. If either a pen or a printer forgets the other, then they simply have to go through the automatic registration procedure again.

When an unknown pen comes within range of a printer, they soon discover they don't know each other. Under these circumstances the pen simply ignores the printer until it finds itself in the charging cup, at which time it initiates the registration procedure.

In addition to its public id, the pen contains a secret id and a secret key-exchange key, both intended for one-time use. These are also recorded in the Netpage Registration Server database at time of manufacture. During registration, the printer obtains the secret id from the pen. Because it is transmitted in the clear, it may be intercepted by someone listening in. The printer transmits the id securely to the Netpage Registration Server, which responds securely with the matching key-exchange key, together with a newly-generated secret id and key-exchange key. The printer generates a session key for the pen and transmits it to the pen encrypted using the one-time-use key-exchange key. It also securely transmits the new secret id and key-exchange key to the pen, which saves them for the next registration procedure. They now match the pen's record in the Netpage Registration Server database.

If the secret id transmitted in the clear from the pen to the printer is intercepted and used to retrieve the secret key-exchange key from the Netpage Registration Server before the printer queries the server, then the server rejects the printer's query because the secret id is out-of-date. Thus the printer knows that the pen has been compromised, and recommends that it be returned for repair.

Whenever a pen is registered, the Netpage Registration Server prints a registration form allowing the pen to be registered in the name of a user. The pen can be registered to an existing user in the registration database, in which case the user's optional password may have to be provided, or new user details can be entered.

The pen uses secret-key rather than public-key encryption because of hardware performance constraints in the pen.

7.4 SECURE DOCUMENTS

The Netpage System supports the delivery of secure documents such as tickets and coupons. The Netpage Printer includes a facility to print watermarks, but will only do so on request from publishers who are suitably authorized. The publisher indicates its authority to print watermarks in its certificate, which the printer is able to authenticate.

The "watermark" printing process uses an alternative dither matrix in specified "watermark" regions of the page. Back-to-back pages contain mirror-image watermark regions which coincide when printed. The dither matrices used in odd and even pages' watermark

regions are designed to produce an interference effect when the regions are viewed together - i.e. when looking *through* the printed sheet.

The effect is similar to a watermark in that it is not visible when looking at only one side of the page, and is lost when the page is copied by normal means.

As described earlier, pages of secure documents cannot be copied using the built-in Netpage copy mechanism. This extends to copying Netpages on Netpage-aware photocopyers.

Secure documents are typically generated as part of e-commerce transactions. They can therefore include the user's photograph which was captured when the user registered biometric information with the Netpage Registration Server, as described in the next section.

7.5 NON-REPUDIATION

In the Netpage System, forms submitted by users are delivered reliably to forms handlers and are persistently archived on Netpage Page Servers. It is therefore impossible for recipients to repudiate delivery.

E-commerce payments made through the system, as described in the next section, are also impossible for the payee to repudiate.

8 Electronic Commerce Model

8.1 SECURE ELECTRONIC TRANSACTION (SET)

The Netpage System uses the Secure Electronic Transaction (SET) [62] system as its payment system model. Although SET is not yet widely supported, it is comprehensive and elegant and will probably become dominant in the near future.

SET, having been developed by MasterCard and Visa, is organized around payment cards, and this is reflected in the terminology. However, much of the system is independent of the type of accounts being used.

In SET, cardholders and merchants register with a certificate authority and are issued with certificates containing their public signature keys. The certificate authority verifies a cardholder's registration details with the card issuer as appropriate, and verifies a merchant's registration details the acquirer as appropriate. Cardholders and merchants store their respective private signature keys securely on their computers. During the payment process, these certificates are used to mutually authenticate a merchant and cardholder, and to authenticate them both to the payment gateway.

SET has not yet been adopted widely, partly because cardholder maintenance of keys and certificates is considered burdensome. Interim solutions which maintain cardholder keys and certificates on a server and give the cardholder access via a password have met with some success [13].

8.2 SET PAYMENTS

In the Netpage System the Netpage Registration Server acts as a proxy for the Netpage user (i.e. the cardholder) in SET payment transactions.

The Netpage System uses biometrics to authenticate the user and authorize SET payments. Because the system is pen-based, the biometric used is the user's on-line signature, consisting of time-varying pen position, tilt and pressure. A fingerprint biometric can also be used by designing a fingerprint sensor into the pen, although at a higher cost. The type of biometric used only affects the capture of the biometric, not the authorization aspects of the system.

The first step to being able to make SET payments is to register the user's biometric with the Netpage Registration Server. This is done in a controlled environment, for example a bank, where the biometric can be captured at the same time as the user's identity is verified. The biometric is captured and stored in the registration database, linked to the user's record and to the record of a particular Netpage Pen. The user's photograph is also optionally captured and linked to the record. The SET cardholder registration process is completed, and the resulting private signature key and certificate are stored in the database. The user's payment card information is also stored, giving the Netpage Registration Server enough information to act as the user's proxy in any SET payment transaction.

When the user eventually supplies the biometric to complete a payment, for example by signing a Netpage order form, the printer securely transmits the order information, the pen id and the biometric data to the Netpage Registration Server. The server verifies the biometric with respect to the user identified by the pen id, and from then on acts as the user's proxy in completing the SET payment transaction.

8.3 MICRO-PAYMENTS

The Netpage Network includes a mechanism for micro-payments, to allow the user to be conveniently charged for printing low-cost documents on demand and for copying copyright documents, and possibly also to allow the user to be reimbursed for expenses incurred in printing advertising material. The latter depends on the level of subsidy already provided to the user.

When the user registers for e-commerce, a network account is established which aggregates micro-payments. The user receives a statement on a regular basis, and can settle any outstanding debit balance using the standard payment mechanism.

The network account can be extended to aggregate subscription fees for periodicals, which would also otherwise be presented to the user in the form of individual statements.

APPLICATIONS AND BUSINESS MODELS

9 Applications

The Netpage Network has the potential to subsume a wide variety of applications in both traditional and electronic media. This section sketches some of the possibilities.

A quick summary of applications:

- personalized newspaper and magazine subscriptions
- freelance columns and bulletin board subscriptions
- on-demand flyers and product brochures
- on-demand books
- e-commerce purchasing from online and traditional catalogs
- delivery of statement and invoices, with online payment facility
- delivery secure document (tickets, coupons and licenses)
- perfect copying with copyright micro-payments
- mail replacement
- delivery of greeting cards
- form printing, fill-in and submission
- delivery of e-mail and facsimile
- Web browsing and printing
- corporate intranets
- government services
- photo album
- computer printing

9.1 PERSONALIZED SUBSCRIPTIONS

The strength of the Netpage Network lies in automatically delivering subscribed periodicals, at a print quality equalling or exceeding that of their traditional counterparts, with editorial content personalized to individual interests, advertising content localized and targeted to individual localities and demographics, and advertising directly linked to detailed product information and product purchasing.

9.1.1 Newspapers and Magazines

The Netpage Network offers a new delivery mechanism to the \$24 billion newspaper and news magazine market which is both more cost-effective than centralized printing and distribution, and allows more fine-grained targeting of advertising.

The simplest form of news personalization involves selecting a news publication and choosing which daily and weekly sections to receive. The simplest form of advertising personalization is tuned to the demographics associated with the subscriber's ZIP or ZIP+4 code. Even with these entry-level forms of personalization, the Netpage Network offers a compelling distribution model for news.

Users receive paper publications with the usability of their traditional counterparts, but with interactive properties. At the press of an on-page button, a user can print an article giving the background to a news story, print a personalized product brochure, or add a product to the virtual shopping basket.

Any magazine normally printed on lightweight paper stock is equally well-suited to distribution on the Netpage Network. However, since Netpage Printers don't carry heavier paper stocks and don't provide wrap-around binding, the Netpage Network is less well-suited to the distribution of so-called "glossy" magazines.

9.1.2 Freelance Columns and Bulletin Boards

Users can choose to subscribe to individual freelance columns, cartoons, etc. These can be integrated into a user's daily news document, or printed individually. Freelancers can choose to receive micro-payments from their subscribers, freeing them from maintaining their own subscriber databases. The Netpage Network provides mechanisms for handling micro-payments.

Users can also subscribe to the "bulletin boards" of friends: collections of news, announcements, pictures etc., which work much like freelance columns.

9.2 ON-DEMAND PUBLICATIONS

The Netpage Network can deliver, on demand, current and back issues of periodicals normally delivered on subscription, including newspapers, magazines, and comics. To maintain the interactivity of all Netpages ever printed, the Netpage Network keeps all published content online at all times. Unlike the Web, where hyperlinks become unreliable over time, content on the Netpage Network never expires.

9.2.1 Flyers and Product Brochures

The Netpage Network makes high-quality flyers and product brochures instantly available, linked to advertisements and entries in printed catalogs.

Brochures are always up-to-date, and link to e-commerce, e-mail, and automatic telephone call-back. Brochure links can provide "click-through" fees to linking documents, and subsidized printing to users.

9.2.2 Books

Users can obtain the latest best-sellers or rare "out-of-print" (a soon-to-obsolete term) editions on demand, printed in column format with a text size chosen by the user. A typical 300-page paperback fits on as little as 40 sheets of Letter paper. Slip-on covers are available for robust handling.

Titles which have outlived their copyright period are available for free. Other titles are heavily discounted for Netpage delivery, since publishers avoid the costs of printing, inventory storage, and delivery.

Colorful children's books reproduce immaculately. When they've been loved to death, they can be printed again, and again.

Children's coloring-in books are available just when they're needed on a rainy day.

9.3 E-COMMERCE

9.3.1 Onlin Purchases

The Netpage Network supports a similar level of online purchasing as the Web, but in paper-based medium which presents like a high-quality printed catalog.

A user can navigate the retailer's online Netpage catalog, printing catalog pages as they're needed and adding items to a virtual shopping trolley. The contents of the trolley can be listed at any time, and items can be struck from the list at the stroke of the pen. Pressing a "proceed to checkout" button at any time elicits a completed order form just waiting for the user's signature. The payment card account number is securely shown in the usual 1234 56** ***** 789 format. The user's signature authorizes the payment.

9.3.2 Catalog Purchases

Rather than buying from an online Netpage catalog, the user can select items from a traditionally-printed catalog which contains active Netpage links.

9.3.3 Statements and Invoices

Statements and invoices can be securely and auditably delivered, and can be automatically filled in with the user's default payment details without the sender knowing those details.

The user's signature can authorize the payment as normal.

9.3.4 Secure Documents

Retailers can securely issue tickets and coupons over the Netpage Network, printed with difficult-to-forge watermarks.

Agencies of various kinds can issue licenses printed with watermarks and the user's own photograph.

9.3.5 Copyright Copying

Any printed version of a Netpage document becomes an easy means to printing another perfect copy. When a copy is made, the Netpage Network can automatically transfer a micro-payment from the copier to the copyright holder.

Trivial copyright fees are universally respected but seldom paid because of the inconvenience. The Netpage Network offers micro-payment convenience and the quality of an original copy.

9.4 COMMUNICATION

9.4.1 Mail

The Netpage Network, once widely subscribed, can be used to deliver numerous instances of regular mail-outs, particularly statements and invoices as discussed above.

9.4.2 Greeting Cards

A user can select a greeting card from an online catalog, add a handwritten message, and dispatch it via the Netpage Network. Cards can be addressed to other Netpage users, and to normal postal addresses. In the latter case the card is printed at the service center closest to the recipient, automatically placed in an envelope, and mailed through the local post.

9.4.3 Forms

Forms of all kinds can be printed on the Netpage Printer, filled in by hand, and submitted directly over the Netpage Network. Submission is secure and cannot be repudiated.

Handwriting is automatically recognized by the system. The digital ink of the handwriting is attached to the form in case a human clerk needs to re-interpret the handwriting. Automatic "handwriting bots" on the network can assist with the recognition task, automatically giving the user semi-intelligent feedback to elicit disambiguation.

Any interactive Netpage "application", including e-commerce and e-mail, uses forms of various kinds.

9.4.4 E-Mail

E-mail forms can be printed on demand and filled in by hand. The handwritten address is converted to facilitate delivery, but the rest of the message is delivered as digital ink, just as the user intended. If the recipient is computer-based rather than Netpage-based, all of the handwriting can be automatically converted, with the digital ink sent as an attachment, since it may contain hand-drawn diagrams etc.

9.4.5 Facsimile

Facsimile forms can be printed on demand and filled in by hand. The handwritten telephone number is converted to facilitate delivery, but the rest of the message is delivered as bitmapped digital ink, just as the user intended.

9.5 WEB BROWSING AND PRINTING

Users can browse the World Wide Web via their Netpage Printer using paper and pen as the user interface. Netpage forms can provide emulation of HTML forms. Only dynamic media objects may fail to print meaningfully.

A Netpage Printer can be the ideal output device for documents encountered while browsing the Web, whether the browsing is terminal-based or Netpage-based. An increasing number of print-ready documents are being published on the Web.

Entire Web sites can be compactly formatted for print, since a printed page has a much greater information-bearing capacity than a computer screen, and perused in a more leisurely fashion.

9.6 CORPORATE INTRANETS

An organization can use a private intranet-based Netpage Network to implement a document repository and efficiently distribute documents on demand.

9.7 GOVERNMENT SERVICES

Government can provide access to services via the Netpage Network. The network can obviate the need to visit government offices to obtain forms and submit forms, and the network can be used to efficiently deliver the results of submissions.

9.8 PHOTO ALBUM

A Netpage Printer can be enhanced with an infrared data connection (IrDA) to allow it to accept images wirelessly from a digital camera. Squirted images can be automatically archived on the Netpage Network, and individual photos, both regular and poster-sized, can be printed at photo quality on the printer.

Smart layout software can assist with the interactive creation of photo album pages, ready for insertion in a cumulative family album.

9.9 COMPUTER PRINTING

A Netpage Printer can be the ideal output device for a personal computer or workstation, whether directly-connected or on a local-area network. The speed of the Netpage Printer and the quality of its output make it attractive to existing computer users.

10 Business Models

10.1 SYSTEM PRINCIPLES

The Netpage Network leverages the open technology and extensive infrastructure of the Internet. The widespread acceptance and growth of the Netpage Network is predicated on open competition rather than monopolistic practices.

However, to provide an incentive to early investors, semi-exclusive licenses to Mem-jet-based Netpage Printer designs will be offered, as well as licenses to manufacture paper and ink consumables.

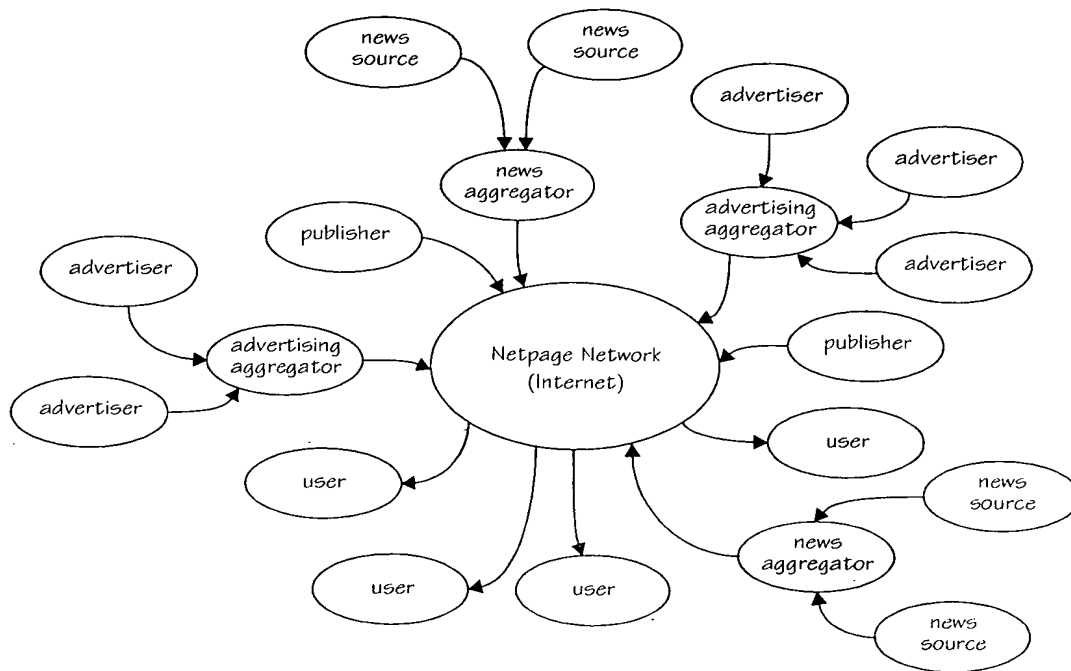


Figure 4. Open structure of Internet-based Netpage Network

The network supports any number of independent participants, some of which have complementary roles, and some of which compete. The open structure of the network is illustrated in Figure 4. Content-related participants include the following:

- news sources
- publishers
- news aggregators
- freelance artists, writers, cartoonists
- direct mailers
- advertisers
- advertising aggregators
- banks
- merchants

Infrastructure-related participants include the following:

- server suppliers
- network storage providers
- communications carriers
- Internet service providers (ISPs)
- printer dealers (plus installation and servicing)
- ink and paper consumables dealers (plus delivery)

Technology-related participants include the following:

- research and development
- chip makers (printheads, controllers, QA)
- printer manufacturers
- ink and paper consumables manufacturers

The strength of the network lies in the fact that publication and delivery are completely decoupled. This allows the delivery infrastructure to grow independently of the participation of publishers.

10.2 BOOTSTRAPPING

Because consumers are unlikely to be motivated to acquire a Netpage Printer until a variety of publications and services are available, and because publishers will wait for an installed base before participating, the key to bootstrapping the network is to bundle the printer with a publication or service subscription, and possibly trimming profit margins in the growth stage.

There are several ways the manufacturing cost of the Netpage Printer, assumed to be well below \$100, can be subsidized. Printer-based distribution can eliminate existing distribution costs, offsetting the printer cost. The printer can provide a new mechanism for delivering advertising, with advertising profits offsetting the printer cost. And the printer cost can be built into the subscription fee for a publication or service.

The cost of printing and delivering a newspaper normally exceeds the price of subscription [31,73]. The real profit lies in the advertising. The cost of the Netpage Printer is easily exceeded by one years' cost savings, allowing a Netpage subscription, including a "free" printer, to be priced lower than a traditional subscription. A Netpage subscription offered to a customer already on the network would be priced correspondingly lower still.

If the publication or service delivered via the Netpage Printer is sufficiently lucrative, then the publisher or provider may be able to subsidize not only the printer itself, but also its running costs. This can include Internet access, paper and ink consumables, and servicing. Demographics-informed advertising may fall into this category. The more information customers reveal about themselves, the greater the value of the advertising to the advertisers, and so the greater the level of subsidization that can take place.

Early investors who subsidize the installation of Netpage Printers may be able to recover the investment and turn a profit merely by charging other publishers an access fee to the printers they "own", perhaps for an interim period after installation, according to a network-wide agreement. They may also be able to earn commissions on click-throughs and e-commerce transactions originating on pages printed on "their" printers.

Similar approaches are already emerging in the general Internet market. In the "FreePC" and related models [49], personal computers are bundled with Internet access, and the whole package is fully or partially subsidized by advertising and e-commerce.

Most content-related participants in the Netpage Network, and even Internet service providers, can benefit from directly investing in Netpage Printer deployment.

10.3 MATURITY

Many of the bundling approaches are likely to remain applicable once the network becomes widely subscribed. It is possible that the bundling of the appliance (i.e. the Netpage Printer) with the service (be it Internet access or a publication subscription) will remain the dominant means of distributing the appliance, as it is in the cellular telephone market, and as an increasing number of companies, including IBM [19], are beginning to believe it should be in the personal computer market.

10.3.1 News Publishers and News Aggregators

News publishers with strong brands are likely to be able leverage those brands on the Netpage Network. They have an incentive to do so quickly to prevent newcomers from filling the vacuum and capturing the attendant advertising revenue.

News publishers also have an incentive to migrate to the Netpage Network because it allows them to offer the more fine-grained targeting that advertisers are increasingly demanding, and which they are increasingly seeking elsewhere.

News publishers who create content rather than simply aggregating other sources have a significant advantage, since they offer both unique content and an editorial voice. Users are more likely to choose a single news publication whose content and editorial orientation they find useful, than specifying to a news aggregator how to glue together a number of disparate news sources. And any sufficiently strong news publication brand is unlikely to make its content available to an aggregator, since the aggregator will be taking a proportion of advertising and e-commerce revenue.

The Netpage Network, like the Web, offers lower barriers to entry than traditional publishing media, and this naturally stimulates greater diversity. However, the geographic independence of the network, coupled with built-in mechanisms for localization of publications, allows international, national and regional news publications to more easily compete in local news markets.

The strength of a traditional local news publication lies partly in its local news content and partly in its local retail advertising and classified advertising content. Aggregation of classified advertising is already happening on the Web, and the Netpage Network will make the same thing possible for local retail advertising. Local news publications are therefore likely to be excluded from the direct capture of local advertising, and may instead transform themselves into news gatherers feeding localized editions of larger publications.

10.3.2 Advertising Aggregators

The Netpage Network promises to be the most effective advertising medium ever conceived. It combines the editorial and print quality of traditional publications with arbitrarily finely targeted advertising, and provides a direct link between advertising, product information, and purchasing.

Because personalization and localization are handled automatically by Netpage Publication Servers, an advertising aggregator can provide arbitrarily broad coverage of both geography and demographics. The subsequent disaggregation is efficient because it is automatic.

This makes it more cost-effective for publishers to deal with advertising aggregators than to directly capture advertising. Even though the advertising aggregator is taking a proportion of advertising revenue, publishers may find the change profit-neutral because of the greater efficiency of aggregation.

Because of the finer targeting supported by the Netpage Network, publishers and advertising aggregators have a larger advertising space to sell, leading to greater profits. The linking between advertising, detailed product information, and purchasing, and the corresponding measurability of consumer behavior, leads to greater profits from click-through fees and e-commerce commissions, benefiting publishers and advertising aggregators alike.

Added revenue from these fees and commissions may even allow users' costs - printer, ink, paper, and Internet access - to be fully subsidized.

NETPAGE PRINTER

11 Printer Product Design

Netpage Printers are intended for use in domestic, commercial, corporate and hospitality environments. They are all based on a straight paper path, passing through a Memjet transfer roller printhead mechanism. In most cases the printed page is glued along one edge and adhered to the previous page to form a final bound document that can be 1 page or 500 pages thick. They all interact with the wireless Netpage Pen.

Netpage Printers are primarily intended for unobtrusive wall-mounting. They also come in tabletop and portable versions.

11.1 WALLPRINTER

A low-cost, wall-mounted, base model with a duplex 8½" Memjet printhead array that accepts a full ream of US Letter paper in a vertical format as shown in Figures 5, 7 and 8. Paper is placed into a hinged top tray down onto a sprung platen and registered under edge guides before being closed. Figure 6 shows the access to the paper and ink cartridge.

A replaceable cartridge containing cyan, magenta, yellow, and infrared inks and glue is also accessible when the tray is open. It connects via a series of self-sealing connectors to hoses that transmit ink and glue to their separate locations. The cartridge consists of a thin wall drawn aluminum casing that accommodates four ink bladders and a single glue bladder into an injection molded connector base. This is a fully recyclable product with a capacity for printing and gluing 3000 pages (1500 sheets). It is protected from forgeries by use of an authentication chip.

When closed, a release mechanism allows the platen to push the paper against the pick-up roller assembly, where it is fed directly into the duplex printhead assembly. From there, the sheet passes a momentary action glue wheel with powered spike wheels, where it has glue applied to the vertical edge as it passes through. The glue wheel is capped when not in use and is operated by a powered camshaft.

The printed sheet is fed down to a binding platen that operates with a closed steel wire loop system of pulleys, runners and a powered axle. This provides the necessary speed to push the sheet forward onto the rear of a previous sheet, glue/bind it and return to the home position to accept the next printed sheet in less than 2 seconds. A motorized paper tapper assembly aligns the sheets in a simultaneous operation.

When a document is bound and finished, a powered exit hatch with a tamper sensor opens. Plastic foils work together with the hatch to direct the finished document to the back of the collection tray and feed any further documents into the tray without hitting existing ones. The collection tray is molded in clear polycarbonate and pulls out of its socket under a certain loading. Access for removing documents is provided on three sides.

The printer has a main PCB that accommodates all major circuitry components including external data jacks. A flex PCB runs from the main PCB to the paper tray and has three different color LEDs and a push button. The LEDs indicate "on", "ink out", "paper out", and "error". The push button is a "help" interface that prints out a simple instruction sheet and a compact features directory for the user. The unit is powered by an internal 110V power supply that is connected before wall-mounting.




Figure 5. Wallprinter




Figure 6. Wallprinter paper and ink cartridge access

The printer has several metal hangers on the rear, which locate into keyhole slots in a metal back plate that is securely fastened to a wall. When the printer has been connected, it is hung onto the back plate and fixed with a locking screw found under the paper tray.

The Netpage Printers are fully customizable in finishes and color as the front moldings clip on to a core chassis and are easily removable.




Figure 7. Wallprinter front elevation





Figure 8. Wallprinter side elevation

11.2 WALLPRINTER PRO

This printer is similar to Wallprinter in most respects, except that it has a duplex 11" print-head assembly, which prints on US Letter paper in a landscape format (see Figures 9, 10 and 11). This means a faster print time and binding time for each page, making for faster overall document delivery.

Another difference is the location of the ink cartridge, which resides above the paper tray rather than down the side. Each page is glued along the horizontal edge by a full-length glue sponge, which is capped when not in use. Operation, printing, and document handling are identical to Wallprinter.

Wallprinter Pro is fully customizable in finishes and color.



Figur 9. Wallprinter Pro




Figure 10. Wallprinter Pro front elevation





Figure 11. Wallprinter Pro side elevation

11.3 TABLEPRINTER PRO

This printer is a tabletop version of the Wallprinter Pro. Essentially, it is the same printer unit with a base plinth that adds extra functionality, such as USB, parallel port and a power socket (see Figure 12).

Tableprinter Pro is fully customizable in finishes and color.



Figur 12. Tabl printer Pro

11.4 WALLPRINTER PRO R

This printer shares the same printing and binding configuration of US Letter landscape format as Wallprinter Pro. The main difference is the media delivery, which is in the form of a large print roll cartridge (see Figures 13, 14 and 15). This cartridge accommodates C, M, Y, and infrared inks and glue as well as a 1000 sheet capacity roll of paper. The cartridge can be recharged at nominated outlets when required and it is protected from forgeries by an authentication chip.

The printer has integral structural metalwork to support its weight and a ball bearing track for easy loading and removal.

Wallprinter Pro R is fully customizable in finishes and color.




Figure 13. Wallprinter Pro R, with print roll extracted




Figure 14. Wallprinter Pro R front elevation




Figure 15. Wallprinter Pro R side elevation

11.5 TABLEPRINTER PRO R

This printer is a tabletop version of the Wallprinter Pro R. Essentially, it is the same printer unit with a base plinth that adds extra functionality, such as USB, parallel port and a power socket (see Figure 16).

Tableprinter Pro R is fully customizable in finishes and color.



Figur 16. Tableprinter Pro R

11.6 TRAVELPRINTER

This is a small, lightweight, versatile and completely portable Netpage Printer (see Figures 17 and 18). It has in-built mobile network communication hardware and software, allowing it to download documents anywhere. Travelprinter also has communications ports for computer interface printing when required.

The printer consists of a front and rear molding with a chassis to accommodate the major components including a lithium battery and an 8½" duplex Memjet printhead assembly.

A compact print roll cartridge with C, M, Y and infrared inks and paper is used in the printer, providing 50 US Letter pages or 100 A5 pages. It is protected from forgeries by an authentication chip.

A motorized guillotine assembly cuts the media between the cartridge and the printhead and motorized spike wheels eject the finished print out of the unit. A flex PCB runs from the main board to a segment LCD and two push buttons. The LCD shows signal strength, any errors, battery and number of pages left in the cartridge. The buttons allow the printer to either connect to the Netpage Network or to act as a stand-alone printer. A USB interface is provided on the side of the printer along with DC 3V input.




Figure 17. Travelprinter, with print roll extracted




Figure 18. Travelprinter R plan and elevations

11.7 NETPAGE PEN

The Netpage Pen is an intelligent, interactive writing implement that communicates with the Netpage Printers and the Netpage Network. It allows the user to write normally with an ink nib or use a non-marking stylus. Rotating the top of the pen selects between them.

The pen has a distinctive shape that is both ergonomic and functional, with imaging optics and electronics housed in the 'underbelly' (see Figures 19, 20 and 21).




Figure 19. Netpage Pen, shown full size

The pen consists of a metal nib that is removable to allow ink cartridge replacement, followed by a tinted diffuser. The diffuser houses an antenna and two LEDs on a flex PCB; the first is located on top of the pen for good all-round visibility. It is a tri-color LED that responds in three modes when the pen is in use: periodic flashing green when it is online

to the printer, momentary green when an operation succeeds, momentary red when an error occurs, yellow when it passes over an active area on the page. A separate lens is mounted into the diffuser, which is optically de-coupled. The lens sits under a second LED that emits infrared light onto the page. The illuminated image is auto-focused onto a controller chip with an on-board image sensor via dedicated optics. The optics chassis accommodates a PCB with various components, including an induction coil for recharging and a MEMS chip. This chip also includes an optical sensor for detecting pressure movement in the metal cam turning when either the stylus or the ink cartridge is used for writing. The cam turning is connected to a terminal collar that has the contact strips for a rechargeable battery. This assembly is fitted into the pen top and connected via two flying leads to the PCB for power transmission. The top assembly is pushed into the body molding where it is free to rotate. By turning the top through 90-degree steps the pen has the stylus out, the pen out or both retracted. The pen has a standard length of 154mm and diameter of 11mm.




Figure 20. Netpage Pen, exploded




Figure 21. Netpage Pen

11.8 PEN CHARGING CUP

The pen charging cup is a simple device that simultaneously recharges any number of Netpage Pens that are placed in it. Figure 22 shows the unit components which are an inner molding, base molding, an induction coil, a termination block, a LED, a diffuser molding and an exit wire to a 3V DC power supply. The inner molding accommodates the induction coil, which is held captive when the unit is assembled. The coil passes through a termination block, where a flying LED is also attached. The LED is positioned at the top of the unit under a diffuser, so it acts as a beacon to indicate that it is active. The charging cup is 70mm in diameter by 93mm high.




Figure 22. Netpage Pen charging cup

12 Memjet-Based Printing

A Memjet printhead produces 1600 dpi bi-level CMYK. On low-diffusion paper, each ejected drop forms an almost perfectly circular 22.5µm diameter dot. Dots are easily produced in isolation, allowing dispersed-dot dithering to be exploited to its fullest. Since the Memjet printhead is the width of the page and operates with a constant paper velocity, the four color planes are printed in perfect registration, allowing ideal dot-on-dot printing. Since there is consequently no spatial interaction between color planes, the same dither matrix is used for each color plane. Dot-on-dot printing minimizes 'muddying' of mid-tones caused by inter-color bleed.

A page layout may contain a mixture of images, graphics and text. Continuous-tone (contone) images and graphics are reproduced using a stochastic dispersed-dot dither. Unlike a clustered-dot (or amplitude-modulated) dither, a *dispersed-dot* (or frequency-modulated) dither reproduces high spatial frequencies (i.e. image detail) almost to the limits of the dot resolution, while simultaneously reproducing lower spatial frequencies to their full color depth, when spatially integrated by the eye. A *stochastic* dither matrix is carefully designed to be free of objectionable low-frequency patterns when tiled across the image. As such its size typically exceeds the minimum size required to support a particular number of intensity levels (e.g. 16×16×8 bits for 257 intensity levels).

Human contrast sensitivity peaks at a spatial frequency of about 3 cycles per degree of visual field and then falls off logarithmically, decreasing by a factor of 100 beyond about 40 cycles per degree and becoming immeasurable beyond 60 cycles per degree. At a normal viewing distance of 12 inches (about 300mm), this translates roughly to 200-300 cycles per inch (cpi) on the printed page, or 400-600 samples per inch according to Nyquist's theorem.

In practice, contone resolution above about 300 ppi is of limited utility outside special applications such as medical imaging. Offset printing of magazines, for example, uses contone resolutions in the range 150 to 300 ppi. Higher resolutions contribute slightly to color error through the dither.

Black text and graphics are reproduced directly using bi-level black dots, and are therefore not antialiased (i.e. low-pass filtered) before being printed. Text is therefore *supersampled* beyond the perceptual limits discussed above, to produce smoother edges when spatially integrated by the eye. Text resolution up to about 1200 dpi continues to contribute to perceived text sharpness (assuming low-diffusion paper, of course).

The Netpage Printer uses a contone resolution of 320 ppi (i.e. 1600÷5), and a black text and graphics resolution of 1600 dpi.

13 Document Data Flow

Because of the high resolution of the Memjet printhead, each page must be printed at a constant speed to avoid creating visible artifacts. This means that the printing speed can't be varied to match the input data rate. Document rasterization and document printing are therefore decoupled to ensure the printhead has a constant supply of data. A page is never printed until it is fully rasterized. This is achieved by storing a compressed version of each rasterized page image in memory.

This decoupling also allows the RIP to run ahead of the printer when rasterizing simple pages, buying time to rasterize more complex pages.

Because contone color images are reproduced by stochastic dithering, but black text and line graphics are reproduced directly using dots, the compressed page image format contains a separate foreground bi-level black layer and background contone color layer. The black layer is composited over the contone layer after the contone layer is dithered.

Figure 23 shows the flow of a Netpage Printer document from network to printed page.

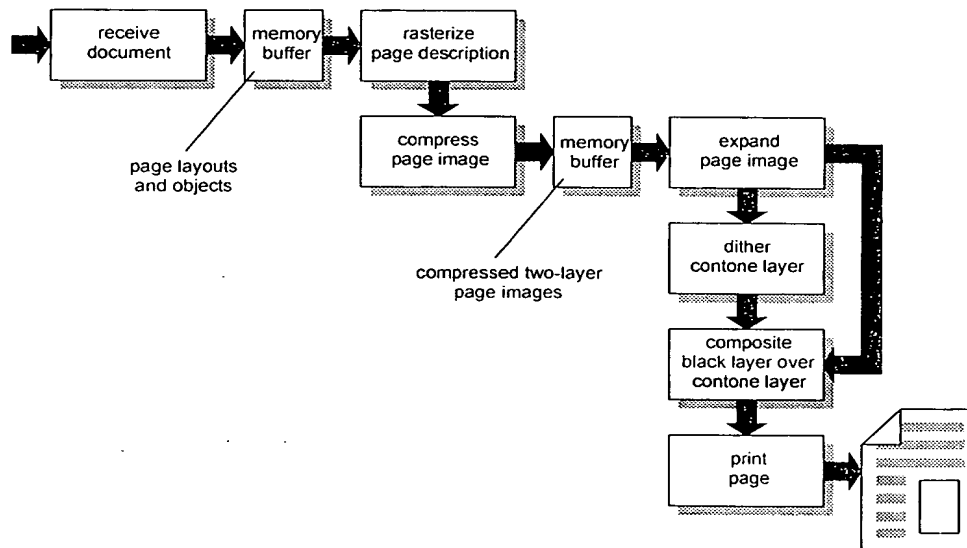


Figure 23. Netpage Printer document data flow

At 320 ppi, a Letter page of contone CMY data has a size of 27MB. Using lossy contone compression algorithms such as JPEG, contone images compress with a ratio up to 10:1 without noticeable loss of quality, giving a compressed page size of 2.7MB.

At 1600 dpi, a Letter page of bi-level data has a size of 29MB. Coherent data such as text compresses very well. Using lossless bi-level compression algorithms such as Group 4 Facsimile, ten-point text compresses with a ratio of about 20:1, giving a compressed page size of 1.4MB.

Once dithered, a page of CMY contone image data consists of 86MB of bi-level data. Using lossless bi-level compression algorithms on this data is pointless precisely because the optimal dither is stochastic - i.e. since it introduces hard-to-compress disorder.

The two-layer compressed page image format therefore exploits the relative strengths of lossy JPEG contone image compression and lossless bi-level text compression. The format is compact enough to be storage-efficient, and simple enough to allow straightforward real-time expansion during printing.

Since text and images normally don't overlap, the normal worst-case page image size is 2.7MB (i.e. image-only), while the normal best-case page image size is 1.4MB (i.e. text-only). The absolute worst-case page image size is 4.1MB (i.e. text over image). Assuming a quarter of an average page contains images, the average page image size is 1.7MB.

14 Printer Controller Architecture

The Netpage Printer controller consists of a controlling processor, a factory-selected network interface, a radio transceiver, dual raster image processor (RIP) DSPs, duplexed print engines, flash memory, and 64MB of DRAM, as illustrated in Figure 24.

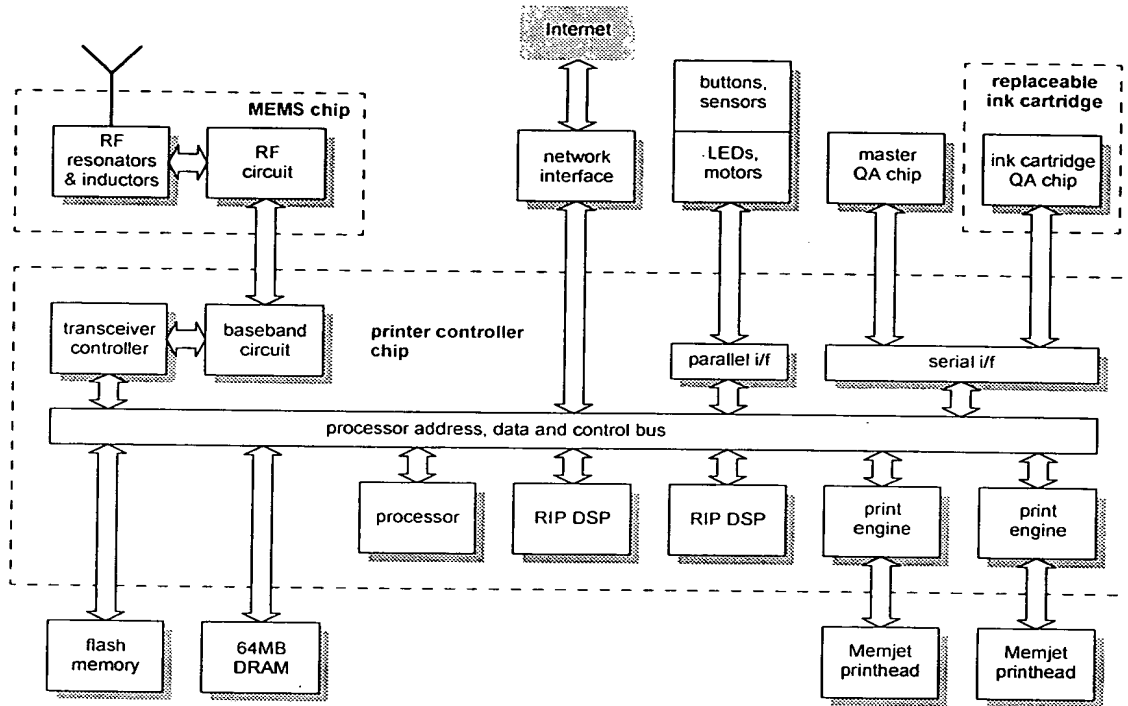


Figure 24. Basic printer controller architecture

The controlling processor handles communication with the Internet and with local wireless pens, controls the user interface (buttons and LEDs), controls the paper transport, handles ink cartridge authentication and ink monitoring, and feeds and synchronizes the RIP DSPs and print engines. It consists of a medium-performance general-purpose microprocessor.

The RIP DSPs rasterize and compress page descriptions to the Netpage Printer's compressed page format. Each print engine expands, dithers and prints page images to its associated Memjet printhead in real time¹. The duplexed print engines print both sides of the page simultaneously.

The printer controller's flash memory holds the software for both the processor and the DSPs, as well as configuration data. This is copied to main memory at boot time.

¹ At 30 or 45 pages per minute.

The processor, DSPs, print engines and digital transceiver components are integrated in a single ASIC. The MEMS and analog RF components are integrated in a separate MEMS chip, which is also used in the wireless pen. Additional pen-specific components in the MEMS chip are not used in the printer controller. The Internet network interface module is separate, since Netpage Printers allow the network connection to be factory-selected. Flash memory and the 2×256Mbit (64MB) DRAM is also off-chip.

Various Internet network interface modules can be supported. Possibilities include a POTS modem, a Hybrid Fiber-Coax (HFC) cable modem, an ISDN modem, a DSL modem, a satellite transceiver, a current or next-generation cellular telephone transceiver, a wireless local loop (WLL) transceiver, etc. A Internet connection may already be available on the local network, in which case only a local network connection may be required.

The printer controller optionally includes a local network connection, to allow the printer to be used from a directly-connected workstation or over a local-area network. Possibilities include 10Base-T and 100Base-T Ethernet, USB and USB 2.0, IEEE 1394 (Firewire), and various emerging home networking standards.

The radio transceiver communicates in the unlicensed 900MHz band normally used by cordless telephones, and uses frequency hopping and collision detection to provide interference-free communication.

14.1 DETAILED DOCUMENT DATA FLOW

The main processor receives and verifies the document's page layouts and page objects by Internet pointcast and multicast as described in earlier sections. It then runs the appropriate RIP software on the DSPs.

The DSPs rasterize each page description and compress the rasterized page image. The main processor stores each compressed page image in memory. The simplest way to load-balance multiple DSPs is to let each DSP rasterize a separate page. The DSPs can always be kept busy since an arbitrary number of rasterized pages can, in general, be stored in memory. This strategy can lead to poor DSP utilization, however, when rasterizing short documents.

Watermark regions in the page description are rasterized to a contone-resolution bi-level bitmap which is losslessly compressed to negligible size and which forms part of the compressed page image.

The infrared (IR) layer of the printed page contains encoded position tags at a density of about 25 per inch. Each tag encodes the page id, tag position, and control bits. Active areas and pressure-sensitive areas in the page description are rasterized to tag-resolution bi-level bitmaps which do not require compression and which form part of the compressed page image.

The main processor passes back-to-back page images to the duplexed print engines. Each print engine stores the compressed page image into its local memory, and starts the page expansion and printing pipeline. Page expansion and printing is pipelined because it is impractical to store a 114MB bi-level CMY+IR page image in memory.

The first stage of the pipeline expands the JPEG-compressed contone CMY layer, expands the Group 4 Fax-compressed bi-level watermark map, and expands the Group 4 Fax-compressed bi-level black layer, all in parallel. The second stage dithers the contone CMY

layer using the dither matrix selected by the watermark map, and composites the bi-level black layer over the resulting bi-level CMY layer. Since there is no black ink used in the Netpage Printer, the black layer is composited with each of C, M and Y. In parallel with this, the tag encoder generates and encodes the bi-level IR tag data. The last stage prints the bi-level CMY+IR data through the Memjet printhead via the printhead interface.

14.2 PRINT ENGINE ARCHITECTURE

The print engine's page expansion and printing pipeline consists of a standard JPEG decoder, a standard Group 4 Fax decoder, a custom ditherer/compositor unit, a custom position tag encoder, and a custom interface to the Memjet printhead. These are described in detail in design documents.

When several print engines are used in unison, such as in a duplexed configuration, they are synchronized via a shared line sync signal. Only one print engine, selected via the external master/slave pin, generates the line sync signal onto the shared line.

In the 8½" versions of the Netpage Printer, each print engine prints 60 Letter pages per minute along the long dimension of the page (11"), giving a line rate of 17.6KHz at 1600 dpi. In the 11" Pro versions of the Netpage Printer, each print engine prints 90 Letter pages per minute along the short dimension of the page (8½"), giving a line rate of 20.4KHz.

These line rates are well within the operating frequency of the Memjet printhead, which in the current design exceeds 30KHz.

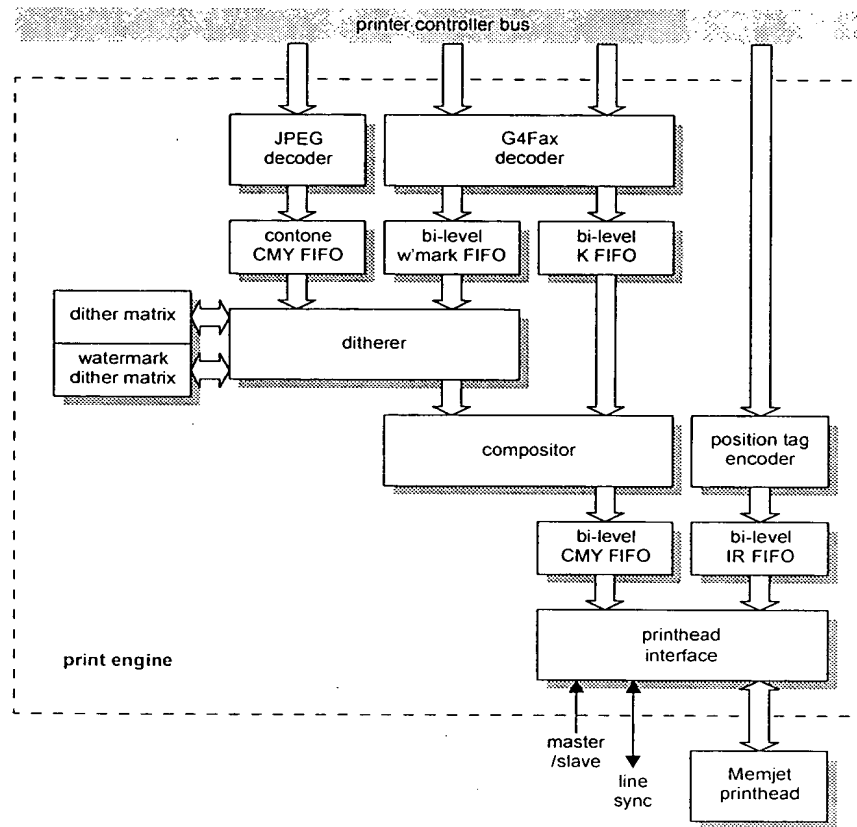


Figure 25. Print engine architecture

15 Pen Controller Architecture

The Netpage Pen controller consists of a controlling processor, a radio transceiver, a tilt sensor, a nib pressure sensor, a IR image sensor, flash memory, and 512KB of DRAM, as illustrated in Figure 26.

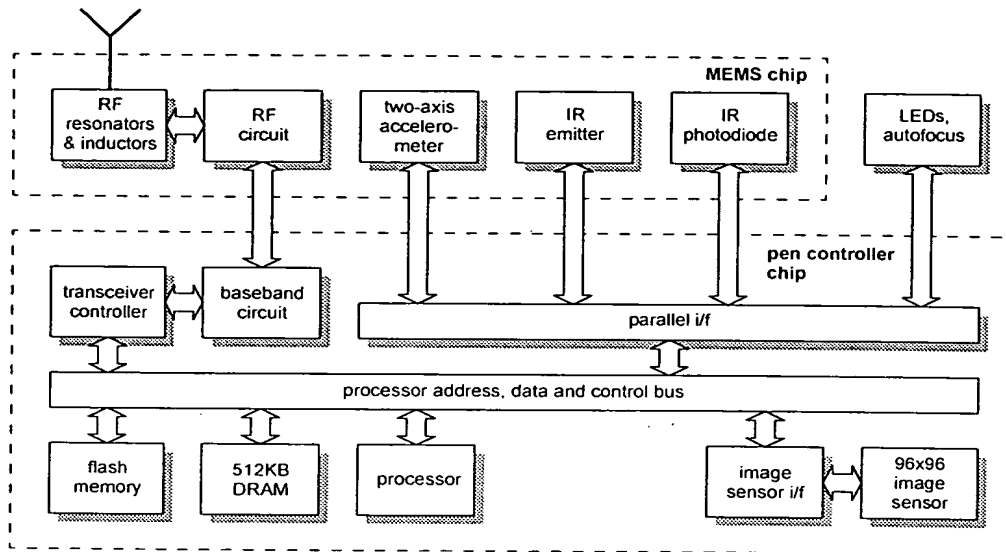


Figure 26. Pen controller architecture

The controlling processor captures and decodes IR position tags from the page via the image sensor, monitors the tilt and pressure sensors, controls the autofocus voice-coil, controls the user interface (tricolor LED), and handles wireless communication with the local Netpage Printer. It consists of a medium-performance (~25MHz RISC) general-purpose microprocessor.

Two-axis tilt sensing is provided by a two-axis accelerometer. Nib pressure sensing is provided by an IR emitter and photodiode pair in conjunction with a reflector coupled with the sprung nib.

The processor, digital transceiver components, 96x96 image sensor, flash memory and 512KB DRAM are integrated in a single ASIC. The MEMS and analog RF components, accelerometers, and the IR emitter/photodiode are integrated in a single MEMS chip, also used in the Netpage Printer.

The radio transceiver communicates in the unlicensed 900MHz band normally used by cordless telephones, and uses frequency hopping and collision detection to provide interference-free communication.

DETAILED DESIGN

16 Position Tag Data Format

The entire page is tiled with position tags. Each position tag must be independent, to allow for single-point virtual button and hyperlink presses, and so contains both the local X-Y position and the global page id. In addition particular flags are associated with each tag, such as whether the pen should turn on its light while over the tag.

The tag data (page id, a local X-Y position, and a number of flags) is error-correctably encoded as a square binary matrix which is printed as an array of bi-level macrodots. Each macrodot itself consists of a square of one or more dots to ease image processing.

The position encoding supports a particular page size. Where the physical page size exceeds this virtual page size, the physical page is simply tiled with multiple virtual pages, each with a different page id. Figure 27 illustrates the concept of a page tiled with position tags. The contents of each tag are not shown on this illustration.




Figure 27. A Page contains a tiling of position tags

The tags are printed in an infrared absorptive ink that can be read by the pen device. Since black ink is IR absorptive, limited functionality can be provided on offset-printed pages using black ink on otherwise blank areas of the page - for example to encode buttons. Alternatively an invisible infrared ink can be used to print the position tags over the top of a regular page. However, if invisible IR ink is used, care must be taken to ensure that any other printed information on the page is printed in infrared-transparent CMY ink, for black ink will obscure the infrared tags. The monochromatic scheme was chosen to maximize dynamic range in blurry reading environments.

Simple virtual button presses don't specifically require invisible ink. Buttons could be visibly coded since the real estate is available. Hyperlinks require invisible ink, but don't require anything else fancy.

Defining the following encoding parameters:

Table 4. Encoding parameters

parameter	definition	typical value
p	virtual page width (inches)	12
d	dot pitch (dots per inch)	1600
k	macrodot width (dots)	4

Table 4. Encoding parameters

parameter	definition	typical value
b	registration border width (macrodots)	2
r	data redundancy factor	1.5
m	coordinate precision (bits)	-
g	page id precision (bits)	-
s	tag width (macrodots)	-
n	tags per page ^a	-
x	position resolution (points)	-

a. per dimension

the tag width is given by:

$$s = b + \sqrt{r(2m + g)} \quad (1)$$

Given a particular tag width, the number of tags per page in each dimension is given by:

$$n = \frac{pd}{ks} \quad (2)$$

This in turn yields the required coordinate precision:

$$m = \log_2 n \quad (3)$$

If 10 billion people all generate 100,000 pages per annum for 1000 years, they will generate 10^{18} pages, or approximately 2^{60} pages. A page id precision of 64 bits should therefore be sufficient, notwithstanding issues of efficient contiguous page id allocation.

Assuming a page id precision of 64 bits, and other parameters as given in Table 4, the equations converge on a tag size of 17 macrodots, 283 tags per page dimension, and a coordinate precision of 9 bits. We round up to 10 bits for future-proofing.

Given a particular tag width, the position resolution is given by:

$$x = \frac{25.4ks}{d} \quad (4)$$

and, *in points*, is given by:

$$x = \frac{72ks}{d} \quad (5)$$

This yields a position resolution of 3 points (1.08mm at 1600 dpi).

Magazine-quality printed text normally has a size of 10 points. Forms filled in by hand normally allow for handwritten text with a size of about 20 points. A position resolution of 3 points therefore translates to a character-oriented resolution of 3.3:1 for printed text and 6.6:1 for handwriting recognition. This position resolution can be increased to the 200 dpi

required for handwriting recognition by taking into account the position of the tag within the captured image area.

With a position tag size of 17 macrodots, and 4 1600 dpi dots per macrodot, the raw position tag gives 24 positions per inch. Consequently any positioning scheme must achieve 10 times better resolution based on tag position within the sensed image, tilt etc.

16.1 POSITION TAG STRUCTURE

Each position tag encodes the following information:

Table 5. Data Encoded in a Position Tag

Name	#Bits	Description
Pageld	64	Defines the Netpage page instance id
X	10	Defines the X coordinate within the given Pageld
Y	10	Defines the Y coordinate within the given Pageld
LightEnable	1	Defines whether the pen should turn its light on while over the tag
ContinuousPressure	1	Determines whether the pen should return continuous pressure or not while over the tag. 0 = don't return pressure readings. 1 = return pressure readings
Reserved	4	For future use. Store as 0 for this version.
TOTAL	90	

Each position tag is a 17×17 array of macrodots, with each macrodot being a 4×4 square of 1600dpi dots, as shown Figure 28.




Figure 28. Macrodots in a Position Tag

Since the dots are monochrome IR absorptive dots are printed on a non-absorptive background a "black macrodot" is physically different from a "white macrodot". Figure 29 illustrates a magnified view of macrodots.




Figure 29. Magnified Macrodots

The position tag structure is based on concepts used in the Aztec 2D Barcode [98] (invented by Andy Longacre of Welch Allyn Inc. in 1995, and is in the public domain). While the position tag uses the same bull's-eye structure mechanism from the Small Aztec Symbol for symbol location and orientation, we define our own interpretation for the mode bits and do not allow variable length structures. In addition, the position tag is a 17×17 array, which is not possible in a pure Aztec code.

Figure 30 shows the high level structure of a position tag in relation to its 17×17 macrodot array. The target/orientation area consists of a 5 level bull's-eye with 3-bit orientation markers on the outer corners. The format of this area is identical to that used in the Small Aztec Barcode although it is placed off-center. The Mode Data Area is the area between the orientation markers in what would have otherwise been the 6th level of the bull's-eye target. The remaining area is used to hold the actual data for the position tag.

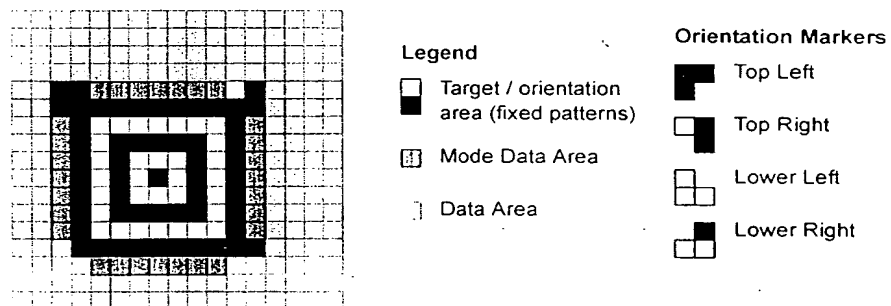


Figure 30. High Level Structure of a Position Tag

16.1.1 Target and Orientation Area

The target and orientation area is a fixed pattern as shown in Figure 31. The format of this area is identical to that used in the Small Aztec Barcode.




Figure 31. Target and Orientation Area

The black positions in are black macrodots, and the white positions are white macrodots. The 5 level bull's-eye pattern is easily found in a 2D image by scans for topological connectivity, then useful for pinpointing the exact center regardless of orientation and for determining the main axes and local X-Y dimensions. On the four corners of the target are 3-bit clusters of orientation bits, which allow the tag's orientation (and even possible mirror imaging) to be quickly determined.

16.1.2 Mode Data Area

The bits in the layer immediately adjoining the target (other than the orientation bits) comprise a 28-bit string, starting upper left and circling clockwise. This area is known as the Mode Data Area. The relative position of the 28 Mode Data bits can be seen in Figure 32.




Figure 32. Mode Data Area

The 28-bit string is encoded meta-information used to decode the remainder of the position tag. Of particular note is the first bit of the meta information. If this bit is 0, then the decoding of the data area is as described here. Setting this bit to 1 allows for redefinition of the position tags at some later date while remaining compatible with existing pages. The meta information and its encoding to 28 bits is described in Section 16.2.1.

16.1.3 Data Area

When the first bit of the meta information decoded from the Mode Data Area is 0, then the data area is defined to be the 168 bits as shown in Figure 30.

The encoded message data is placed onto the tag in 2-cell high layers, starting at the upper left and spiraling clockwise out to the edge of the symbol. The data is therefore interpreted as a sequence of related "dominos", each 1 wide and 2 tall, with their more significant bit

always further from the target area (a black macrodot represents a “1” bit, and a white macrodot represents a “0” bit). Figure 30 illustrates the sequence and orientation of the dominos when turning the corners in the data area.




Figure 33. Domino Mapping and Sequencing of Data Area

The systematic domino-based layout of the data simplifies both encoding and decoding at the graphical level. The encoding mechanism for the data is described in Section 16.2.2, although it is worth noting here that decoding errors and tag damage are expected to occur more towards the edges of the tag.

Note that the 17×17 structure of the position tag means that with reference to the strict Aztec Code, the first layer of data is complete, but the second layer of data is only half full. If the second layer was completely full the position tag would be a 19×19 square cell with the target area in the center.

16.2 POSITION TAG ENCODING

16.2.1 The Mode Message

The 28-bit string in the Mode Area is used to hold 8 mode bits. These 8 mode bits describe how the remainder of the position tag bits were encoded.

The 8 mode bits are encoded into 28 bits via Reed Solomon Encoding. The Mode Area therefore holds the 8 mode bits and 20 additional check bits. The 8 mode bits are parsed into two 4-bit words, and then 5 additional check words are computed by systematic Reed-Solomon encoding over the Galois field GF(16) based on a prime modulus polynomial of: $x^4 + x + 1$ (=19 decimal). The generator polynomial of $(x-2^1) \dots (x-2^4)$ is:

$$x^4 + 11x^3 + 4x^2 + 6x + 2$$

Reed-Solomon encoding was chosen for its ability to deal with burst errors and effectively detect and correct errors using a minimum of redundancy. Reed Solomon encoding is adequately discussed in [99], [82], and [55]. The reader is advised to refer to these sources for background information.

Table 6 shows the interpretation of the mode bits. Note that the first bit of the mode data is a version bit, and determines the interpretation of the remainder of the mode bits and the

tag data. The Version bit should therefore be set to 0. Later interpretations of the tag can be defined at a later date via the version bit.

Table 6. Interpretation of Mode Bits

#Bits	Name	Description
0	Version	0 = the interpretation of the remaining 7 bits and the data area is as described in this document 1 = for future use
1-2	NumMsgWords	Defines how many 6-bit codewords of the data area are message codewords, and how many are check-word codewords. This number is determined by the encoding scheme used for data (see Section 16.2.2). Values are: 00 = 15 data message codewords 01 = 16 data message codewords 10 = 17 data message codewords 11 = 18 data message codewords
3-7	Reserved	0

16.2.2 The Data Message

The Data Area contains 168 raw bits, which is an encoded form of the 90 bits of position tag data (see Table 5). The 168 bits are represented by 104 bits in the first layer surrounding the target area and 64 bits in the layer half surrounding the first.

The 90 bits of position tag data are encoded into 168 bits using two steps. The first step is to break the 90 bits into codewords of 6 bits such that there are no codewords with all 0s or all 1s. The second step is to Reed-Solomon encode the data over the Galois field GF(64), which generates additional 6-bit check words.

16.2.2.1 Step 1: Message Encoding

In the first step of message encoding, the 90 data bits is laid into a sequence of 6-bit message codewords in a generally direct fashion, starting at the most significant bit of the first codeword, with two key exceptions:

- Whenever the first 5 bits placed in a codeword are all “0”s, then a dummy “1” is inserted into that codeword’s LSB and the following message bit starts off the next codeword.
- Whenever the first 5 bits placed in a codeword are all “1”s, then a dummy “0” is inserted into that codeword’s LSB and the following message bit starts off the next codeword.

In the end, the character and byte boundaries in the original 90-bit message have no necessary relationship with the codeword boundaries. Up to 5 bits may remain unfilled in the final message codeword, and they are to be padded out with “1”s (and possibly a final dummy “0” if necessary) to eliminate any ambiguity.

The code-forming rules are designed to never create a *message* codeword of all 0s or all 1s, but the error encoding adds on additional codewords of any value. Thus during decoding, an occurrence of those illegal values within the message region (but *not* within the check region) can be regarded as a correctable *erasure* (this is more useful than not knowing which codeword is in error).

In the case of our 90-bit original messages, there are two cases to consider:

- the *best* case is that 15 code words are generated for the data message
- the *worst* case is that 18 code words are generated for the data message

The codeword based message therefore consumes D 6-bit codewords, where D is in the range 15-18.

16.2.2.2 Step 2: Generate Check Words

The second step of message encoding involves generating the check words for the data message. Since the data message consumes between 15 and 18 6-bit of the 28 available codewords, the remaining 13 to 10 codewords can be used for check words.

The K additional check words are computed by systematic Reed-Solomon encoding over the Galois field GF(64) based on a prime modulus polynomial of: $x^6 + x + 1$ (=67 decimal). The generator polynomial is simply expanded out to $(x-2^1) \dots (x-2^K)$ as needed at the time of printing. The check words are not adjusted in the same way as the message bits are adjusted.

16.2.2.3 Step 3: Place into Data Area

The final message to be stored in the data area of a position tag consists of 28 6-bit codewords made up of D message codewords followed by K check codewords. The exact number of message codewords, D , is stored into the 2-bit NumMsgWords component of the Mode information for the position tag.

The resulting sequence of codewords is parsed anew into a sequence of 2-bit dominos (see Figure 34) and is then taken in reverse order and graphically laid spiralling clockwise and outward through the data area (see Figure 33). As decoding errors and tag damage are expected to occur more towards the edges of the tag, by reversing the codeword sequence, the message codewords occupy those edge regions.




Figure 34. Relationship of 6-bit codeword to generated dominos

One macrodot is placed into the data area per bit. The placement of an IR absorptive ink macrodot represents a "1". If there is no IR absorptive ink placed on the page at the particular location, then that macrodot represents a "0".

17 In-Pen Processing

The pen must send the appliance 100 samples per second. The mechanism for reading positional information from the position tags must work at this 100 Hz rate. The following tasks must be performed each sample:

- Capture time, tilt and pressure
- Capture the image
- Locate the position tag
- Decode the position tag
- Build stroke information
- Activate the light as required
- Encrypt the stroke data
- Transmit the encrypted data

When a pen down is registered, we capture the first three frames at the two extreme focus values and the middle focus value. After that we perform regular auto focus based on a sharpness metric. In this way we are more likely to capture a “single-click” event.

While the pen is in range of the receiver unit, partially completed strokes are transmitted as they are being formed. When the pen moves out of range, stroke information is buffered within the pen (approximately 12 minutes of pen motion while physically on the page), and transmitted later.

Each pen contains a characteristics block which can be read by the appliance. It contains at least the information shown in Table 7:

Table 7. Pen Characteristics Block

Name	Description
PenId	A unique pen identifier - unique across all pens
ModelId	The model number of this pen ^a
ManufacturerId	The manufacturer of this pen
TiltInfo	The relationship between tilt and pen position.

- a. This can be combined with ManufacturerId to allow the appliance to download a set of pen characteristics. However it may be more convenient to store the characteristics (such as TiltInfo) in the pen for convenience.

17.1 POSITION TAG DECODING

17.1.1 Capture Time, Tilt and Pressure

The pen contains a number of sensors, including current pressure and tilt.

17.1.1.1 Time

All measurements are made in the context of a particular time stamp. The pen contains a timer to enable equally time-spaced measurements to be made. Although we capture samples at a rate of 100 per second, we report time at the millisecond resolution to allow for future improvements to the pen. Time resolution is provided at 32 bits.

17.1.1.2 Tilt

The pen contains 2 tilt sensors - each one directly measures tilt for a given axis with a resolution accuracy to at least 5 bits.

Tilt is required to determine the nib position. The appliance uses the two tilt values and the pen's non-corrected pen position to determine the actual pen nib position. The Pen Characteristics Block contains the relationship required for the calculation. Offloading the correction calculation to the appliance reduces pen complexity and price.

17.1.1.3 Pressure

The pen contains a pressure sensor which measures to an accuracy of at least 5 bits. From the pen's point of view, the pressure value is only used to determine whether the pen is on the paper or not. The appliance is responsible for all other interpretations of pressure, such as determining whether an item on a page is being selected. The pen merely passes the pressure value on to the appliance.

Each position tag contains a data bit called ContinuousPressure. If a given pen stroke includes position tags that have their ContinuousPressure bit set, then the pressure value is passed as part of the stroke information. If the ContinuousPressure is clear, then pressure is implied from a stroke in terms of pen-down and pen-up.

17.1.2 Capture the Image

The image needs to be captured for the pen position to be determined. This only needs to happen if the pen is on the page. If the pen is not on the page then capturing an image will be a waste of time.

To capture an image, data must be transferred from the image sensor to a fixed image buffer. This process involves a single read and write per pixel. The minimum and maximum values encountered during the transfer are kept in variables MinPixelEncountered and MaxPixelEncountered respectively for later use in the data recovery process (see Section 17.1.4 on page 92).

There are a number of general considerations that are part of the assumptions for reading in and decoding a position tag from an image sensor.

With regards to the image sensor itself, there are two calculations to consider:

- sensed area calculation
- image sensor calculation

In summary, the CMOS image sensor must be able to sense 96×96 pixels over a 3.05mm × 3.05mm area. This implies an 800dpi sensor.

Finally, the sensed image will be blurry.

A 96×96 image contains 9,216 8-bit pixels. We rounding up to 10,000 pixels for convenience. The transfer process involves on average 6.5 cycles per pixel for a total of 65,000 cycles per position sample. On a 25 MHz processor this represents 26% of the available bandwidth.

17.1.2.1 *Sensed Area*

The basic positional tag is 17×17 macrodots. Since each macrodot is a 4×4 array of printed dots, each macrodot is in fact 68×68 printed dots.

To satisfy Nyquist's Theorem, the sensor must be able to see at least 2 tags. If the sensor is only large enough to see a single tag, it is most likely to see part of one tag and part of another. By being wide enough to see two tags we guarantee being able to see at least one complete tag. This requirement increases the sensed dot area to 136×136 .

Finally, since the pen must detect positional tags at any rotation, we must consider the worst case rotation of 45 degrees. This requirement increases the sensed dot area to 192×192 .

At a printed dot size of 1600 dpi, the image sensor must sense a printed area containing 192×192 dots, which equates to 3.05mm (0.12 inches) in each dimension.

17.1.2.2 *Image Sensor Resolution*

The basic positional tag is 17×17 macrodots. This equates to a basic sensor resolution of 17×17 pixels.

To satisfy Nyquist's Theorem with respect to the positional tags, the sensor must be able to see at least 2 tags. If the sensor is only large enough to see a single tag, it is most likely to see part of one tag and part of another. By being wide enough to see two tags we guarantee being able to see at least one complete tag. This requirement increases the sensor resolution to 34×34 .

Since the pen must detect positional tags at any rotation, we must consider the worst case rotation of 45 degrees. This requirement increases the sensor resolution to 48×48 .

Finally, to satisfy Nyquist's Theorem with respect to macrodots, we must oversample at least at double the macrodot resolution. This leads to a sensor resolution of 96×96 .

Each sampled pixel is 1 byte (8 bits). We must also assume the lowest 2 bits of each pixel can contain noise. Decoding algorithms must therefore be noise tolerant.

17.1.2.3 *Blurry Image*

Blurring is introduced into the sensed image in the following ways:

- Natural blurring due to nature of the sensor's distance from the page (focus)
- Variable blurring over the image due to pen tilt
- Warping of the page

Natural blurring of an image occurs because it is sensed out of focus. Consequently there is overlap of sensed data in the area sensors. Blurring is useful, as the overlap ensures there are no high frequencies in the sensed data, and that there is no data missed by the sensor. However if the area covered by a sensor pixel is too large, there will be too much blurring and the sampling required to recover the data will not be met.

The second form of blurring occurs because the pen has a varied tilt. Because of the pen's tilt, part of the image will be in focus, while another part will be slightly out of focus. The more pronounced the tilt the greater the focus difference from one part of the image and

another. The sensor lens is specifically designed to minimize the differences in blurring across the image, but it will still be there.

The other form of blurring occurs when a page is slightly warped. When the warping is in the vertical dimension, the distance between the page and the sensor will not be constant, and the level of blurring will vary across those areas.

Monochrome dots (IR reflective and non-IR reflective) were chosen for the tags to give the best dynamic range in blurry reading environments. Blurring can cause problems in attempting to determine whether a given dot is “black” or “white”. Figure 35 shows the effect of blurring on a position tag composed of macrodots, from no blurring to severe blurring. Notice how isolated macrodots disappear as the blurring becomes more severe.




Figure 35. Effect of Different Levels of Blurring

As the blurring increases, the more a given macrodot is influenced by the surrounding macrodots. Consequently the dynamic range for a particular macrodot decreases. Consider a “white” macrodot and a “black” macrodot, each surrounded by all possible sets of macrodots. The 9 macrodots are blurred, and the center macrodot sampled. Figure 36 shows the distribution of resultant center macrodot pixel values for black and white macrodots.




Figure 36. Generalized blurring curve

Figure 36 is intended to be a representative blurring. The curve from pixel values 0-180 shows the range of black macrodots. The curve from pixel values 75-250 shows the range of white macrodots. However the greater the blurring, the more the two curves shift towards the center of the range and therefore the greater the intersection area, which means the more difficult it is to determine whether a given macrodot is black or white. A pixel value at the center point of intersection is ambiguous - the macrodot is equally likely to be black or white.

As the blurring increases, the likelihood of a read bit error increases. Fortunately, the Reed-Solomon decoding algorithm can cope with these gracefully up to t symbol errors. Figure 37 shows the predicted number of Reed-Solomon codewords that cannot be recovered given a particular symbol error rate. Notice how the Reed-Solomon decoding scheme performs well and then substantially degrades. If there is no Reed-Solomon block duplication, then only 1 codeword needs to be in error for the data to be unrecoverable. Of course, with the inclusion of erasure detection (via invalid codewords), the chance of correctly decoding the data increases.




Figure 37. Relationship between data non-recovery and codeword error rate

Figure 37 only links codeword errors to the total number of Reed-Solomon codewords in error. There is a trade-off between the amount of blurring that can be coped with, compared to the amount of damage that has been done to a card. Since all error detection and correction is performed by a Reed-Solomon decoder, there is a finite number of errors per Reed-Solomon data block that can be coped with. The more errors introduced through blurring, the fewer the number of errors that can be coped with due to page damage.

17.1.3 Locate the Position Tag

The captured image is a 96x96 pixel map of 8-bit samples. Before any position information can be decoded, the position tag must be located within the image. Note that the tag can have any orientation.

The first step in locating the tag is to locate a bull's-eye target (Figure 31) within the image. Since we are imaging a wider area than a single tag it is possible that two targets lie within the image area.

Since an entire tag must be completely contained within the image, the data area outside the target must be visible. As a result, the area of interest for locating targets is actually a

subset of the whole image. Instead of searching a 96×96 image, we only need search an 88×88 image (a reduction of 15% from 9216 samples to 7744).




Figure 38. Area of Interest for Target Location

The 5 level bull's-eye pattern of the target is easily found in a 2D image by scans for topological connectivity, then useful for pinpointing the exact center regardless of orientation and for determining the main axes and local X-Y dimensions.

17.1.3.1 Scan for Topological Connectivity

Prior to image traversal we initialize an 88×88 array of flags with 0. This array keeps the visited status of the corresponding pixel in the image. If the flag is 0 we have not yet visited the pixel. If it is 1 or 2, then we have visited the pixel and bit 1 of the flag contains the sense of the bit (1 for black, 0 for white).

We then start with the top left pixel and compare adjacent pixels in order to build a connectivity structure for the top left pixel. If the top left pixel is considered to be white, then the connectivity structure will represent all connected pixels that are also white. If the top pixel is considered to be black, then the connectivity structure will represent all connected pixels that are also black. As edges are met due to differences in color sense, adjacent connectivity structures are built.

The order of pixel traversal in the image simply follows a standard pixel fill algorithm. There are many such algorithms to be found since this is a useful operation for paint programs. Near-optimal algorithms such as [29] and [37] only read each pixel once on average.

In our case we must also consider the notion of color sense. Target location does not require the sophistication of the macrodot reading algorithm for mode bits and data bits for determining if a macrodot is black or white (Section 17.1.4.2 and Section 17.1.4.4). The minimum and maximum pixels encountered during the image capture were stored in `MinPixelEncountered` and `MaxPixelEncountered` respectively, and these are used to set the threshold for black/white as $\text{MidRange} = (\text{MinPixelEncountered} + \text{MaxPixelEncountered})/2$. Anything greater than or equal to `MidRange` is considered to be white, and anything less than `MidRange` is considered to be black.

As each connectivity structures is completed, it is compared against the expected values for the 5 bull's-eye layers surrounding the target center. Possible matches are added to one list, while non-matches are discarded.

The total time taken for processing the image is approximately 6 cycles per sample pixel. This equates to 46,464 cycles (88×88×6), or 19% of the available processor bandwidth.

17.1.3.2 Determine Target Axes

The connectivity structures can now be examined together in order to locate the target. The identification of 2 or more concentric structures is enough to determine the X/Y axes.

Taking one of the completely isolated concentric structures, 2 points on each of the 4 sides are chosen where the distance separating them is at least 3 pixels. For each point, the pixel of maximum whiteness or maximum blackness is chosen as the estimate of the center of the line. The estimate should be within 1 pixel of the actual center.

The process of building a more accurate position for the line center involves reconstructing the continuous signal for 7 scanline slices of the line, 3 to either side of the estimated center. The 7 maximum values found (one for each of these pixel dimension slices) are then used to reconstruct a continuous signal in the column dimension and thus to locate the maximum value in that dimension.

```
// Given estimates column and pixel, determine a
// betterColumn and betterPixel as the center of the target
for (y=0; y<7; y++)
{
    for (x=0; x<7; x++)
        samples[x] = GetPixel(column-3+y, pixel-3+x)
    FindMax(samples, pos, maxVal)
    reSamples[y] = maxVal
    if (y == 3)
        betterPixel = pos + pixel
}
FindMax(reSamples, pos, maxVal)
betterColumn = pos + column
```

FindMax is a function that reconstructs the original 1 dimensional signal based sample points and returns the position of the maximum as well as the maximum value found. The method of signal reconstruction/resampling used is the Lanczos3 windowed sinc function (see pages 157/158 in [96]). The function and kernel are shown in Figure 39:




Figure 39. Lanczos3 windowed sinc function

The Lanczos3 windowed sinc function takes 7 (pixel) samples from the dimension being reconstructed, centered around the estimated position X , i.e. at $X-3$, $X-2$, $X-1$, X , $X+1$, $X+2$, $X+3$. We reconstruct points from $X-1$ to $X+1$, each at an interval of 0.1, and determine which point is the maximum. The position that is the maximum value becomes the new center. Due to the nature of the kernel, we only require 6 entries in the convolution

kernel for points between X and $X+1$. We use 6 points for $X-1$ to X , and 6 points for X to $X+1$, requiring 7 points overall in order to get pixel values from $X-1$ to $X+1$ since some of the pixels required are the same.

Two accurate positions on each of 2 sides of the isolated bull's-eye layer give the X and Y axes directly. Using 4 sides allows an average line to be calculated (simply the average of the two lines built from opposite sides).

The distance from the center of one bull's-eye layer macrodot line to the center of the next bull's-eye layer's macrodot line gives the distance between macrodots in terms of a ΔX and a ΔY with respect to the current orientation of the image.

17.1.3.3 Determine Target Center

In order to locate the actual center of the target there are two cases to consider.

- The center of the target is one of the connectivity structures
- The center of the target is *not* one of the connectivity structures

In the first case, we choose the blackest pixel of the estimated target center, which should be within 1 pixel of the actual center. The process of building a more accurate position for the target center involves reconstructing the continuous signal for 7 scanline slices of the target, 3 to either side of the estimated target center. The 7 maximum values found (one for each of these pixel dimension slices) are then used to reconstruct a continuous signal in the column dimension and thus to locate the maximum value in that dimension. The process is the same as described in Section 17.1.3.2 and uses the same filter as shown in Figure 39.

If there is no current estimation for the center, the center is considered to be the calculated center of the four lines used in the calculation of the axes in Section 17.1.3.2. This process can be repeated for another isolated bull's-eye layer and the two estimates for target center averaged. If the estimated center pixel is black, a more accurate center can potentially be obtained by following the procedure defined for the first case. However care must be taken since this pixel region is not completely isolated (or there would have been a target estimate in the connectivity structures). It may be enough to increase the accuracy in only one dimension.

17.1.4 Decode the Position Tag

With the position of the tag known, the task of decoding it can begin. Decoding a positional tag consists of 3 essential steps:

- Determining the orientation of the tag
- Extracting the tag's mode data and decoding it
- Extracting and decoding the data portion of the tag based on the mode data

17.1.4.1 Determine Orientation

On the four corners of the target are 3-bit clusters of orientation bits, which allow the tag's orientation (and even possible mirror imaging) to be quickly determined.

Using the ΔX and ΔY values (obtained in Section 17.1.3.2), and applying them to the accurate target center coordinate (obtained in Section 17.1.3.3), we are able to calculate the estimated positions of the 12 orientation macrodots.

Given the macrodot coordinate (fixed point) we sample 4 image pixels to arrive at a center pixel value via bilinear interpolation.

Once the center pixel value has been determined, we try to determine the bit value for that macrodot. To do so, we take the pixel values representing the macrodot centers to either side of the macrodot whose bit value is being determined, and attempt to intelligently guess the value of that center macrodot's bit value. Looking at the generalized blurring curve again (reproduced in Figure 40 from Figure 36), but with BlackMax and WhiteMin shown, there are three common cases to consider:




Figure 40. Generalized blurring curve

- The macrodot's center pixel value is lower than WhiteMin, and is therefore definitely a black macrodot. The bit value is therefore definitely 1.
- The macrodot's center pixel value is higher than BlackMax, and is therefore definitely a white macrodot. The bit value is therefore definitely 0.
- The macrodot's center pixel value is somewhere between BlackMax and WhiteMin. The macrodot may be black, and it may be white. The value for the bit is therefore in question. A number of schemes can be devised to make a reasonable guess as to the value of the bit. These schemes must balance complexity against accuracy, and also take into account the fact that in some cases, there is no guaranteed solution. In those cases where we make a wrong bit decision, the bit will be in error, and must be corrected by some other means (in the case of orientation bits, the other corners may help, and in the case of the tag's data, the Reed-Solomon decoding stage will correct errors).

Before the data can be extracted from the data area, the pixel ranges for "black" and white dots needs to be ascertained. The minimum and maximum pixels encountered during the search for the tag were stored in MinPixelEncountered and MaxPixelEncountered respectively. The following pseudocode shows the method of obtaining good values for WhiteMin and BlackMax based on the min & max pixels encountered.

```

MinPixel = MinPixelEncountered
MaxPixel = MaxPixelEncountered
MidRange = (MinPixel + MaxPixel) / 2
WhiteMin = MaxPixel - PenConstantWhiteFactor
BlackMax = MinPixel + PenConstantBlackFactor

```

The scheme used to determine a macrodot's value if the pixel value is between BlackMax and WhiteMin is fairly simple, but gives good results. It uses the pixel values of the macro-

dot centers to the left and right of the dot in question, using their values to help determine a more likely value for the center dot:

- If the two macrodots to either side are on the white side of MidRange (an average macrodot value), then we can guess that if the center macrodot were white, it would likely be a “definite” white. The fact that it is in the not-sure region would indicate that the macrodot was black, and had been affected by the surrounding white macrodots to make the value less sure. The macrodot value is therefore assumed to be black, and hence the bit value is 1.
- If the two macrodots to either side are on the black side of MidRange, then we can guess that if the center macrodot were black, it would likely be a “definite” black. The fact that it is in the not-sure region would indicate that the macrodot was white, and had been affected by the surrounding black macrodots to make the value less sure. The macrodot value is therefore assumed to be white, and hence the bit value is 0.
- If one macrodot is on the black side of MidRange, and the other macrodot is on the white side of MidRange, we simply use the center macrodot value to decide. If the center macrodot is on the black side of MidRange, we choose black (bit value 1). Otherwise we choose white (bit value 0).

The logic is represented by this simple pseudocode:

```

if (pixel < WhiteMin)           // definitely black
    bit = 0x01
else
if (pixel > BlackMax)           // definitely white
    bit = 0x00
else
if ((prev > MidRange) && (next > MidRange)) //prob black
    bit = 0x01
else
if ((prev < MidRange) && (next < MidRange)) //prob white
    bit = 0x00
else
if (pixel < MidRange)
    bit = 0x01
else
    bit = 0x00

```

From this one can see that using surrounding pixel values can give a good indication of the value of the center macrodot's state. The scheme described here only uses the macrodots from the same row, but using a single macrodot line history (the previous macrodot line) would be straightforward and worth investigating, to see if the added complexity of using additional neighboring macrodots would give a useful increase in accuracy. This is an interesting area for future research.

We already have the distance between the macrodots in terms of a ΔX and a ΔY with respect to the current orientation of the image. Rather than rotate the image to transform ΔX and ΔY both to 1, we simply change the sense of the ΔX and ΔY values based on how many sets of 90 degrees it takes for the tag to be oriented so that both D values are positive. The four possible rotations by 90 degrees equate to 4 different combinations of signs ($\pm\Delta X$ and $\pm\Delta Y$). It is therefore trivial to translate ΔX and ΔY into two deltas for movement in X and Y: ΔX_X and ΔY_X for movement of 1 macrodot in X, and ΔX_Y and ΔY_Y for movement of 1 macrodot in Y.

With the 12 orientation bits obtained, the orientation of the tag can be easily determined. Given 4 sets of 12 bits, the correct orientation is the one that minimizes the number of bits in error, as shown in Table 8.

Table 8. The Four Major Orientations of the Position Tag

12-bit value	Orientation of tag
111-011-100-000	No rotation
011-100-000-111	Rotated clockwise 90 degrees
100-000-111-011	Rotated clockwise 180 degrees
000-111-011-100	Rotated clockwise 270 degrees

17.1.4.2 Extract the Tag's Mode Raw Data

With the orientation determined, the next step is to read the bit pattern representing the Mode bits from the tag.

We know the coordinate of the tag's target center as $[Center_X, Center_Y]$. We also know the orientation of the tag and therefore have ΔX_X and ΔY_X for motion in X, and ΔX_Y and ΔY_Y for movement in Y. These delta values enable us to step along the center of any macrodots. By sampling the tag at the 18 mode bit positions (see Figure 32). Consequently the location of the first few mode bit are:

- Bit0 location = $[Center_X - 5\Delta X_X - 5\Delta X_Y, Center_Y - 5\Delta Y_X - 5\Delta Y_Y]$
- Bit1 location = $[Bit0_X + \Delta X_X, Bit0_Y + \Delta Y_X]$
- Bit2 location = $[Bit1_X + \Delta X_X, Bit1_Y + \Delta Y_X]$

Given the macrodot coordinate (fixed point) we use the sampling mechanism as described in Section 17.1.4.1 to extract the bit value.

The extraction of 28 bits will take approximately 168 cycles on a simple microprocessor. Assuming a clock speed of 25 MHz, 168 cycles per position sample represents less than 0.1% of the available processor bandwidth.

17.1.4.3 Decode the Mode Data

The 28-bit raw bitstream representing the tag's mode data has been read in. It must now be decoded into 8 bits of meaningful mode information.

The data is decoded using straightforward Reed-Solomon decoding. One of the advantages of using Galois Field GF(16) is that a total of only 64 bytes are required for lookup tables during the decoding process (16 bytes each of power, log, multiply and multiply inverse).

The decoding process of 28 bits to 8 bits will take approximately 1,500 cycles on a simple microprocessor. Assuming a clock speed of 25 MHz, 1,500 cycles per position sample is less than 1% of the available processor bandwidth.

17.1.4.4 Extract the Raw Tag Data

Before the tag data *proper* can be read in from the tag image, the tag's 8-bit mode data must be examined. The Version bit must be checked to ensure that this tag has a known structure (see Table 6). If the Version bit is 0, then the tag data structure is known and can be read in from the image.

The macrodots representing the tag data area are read into a 168-bit raw bitstream according to the locations described in Figure 33. We use the tag's target center [$Center_x$, $Center_y$] and the orientation delta values ΔX_x and ΔY_x , ΔX_y and ΔY_y to calculate the start position. The delta values also enable us to step along the center of the data area's macrodots.

The process of determining if a given macrodot is a 1 or a 0 is the same as that carried out for reading the tag's mode bits and orientation bits (see Section 17.1.4.1). A total of 168 bits is read in from the image in this way. The extraction of 168 bits will take approximately 1008 cycles on a simple microprocessor. Assuming a clock speed of 25 MHz, 1008 cycles per position sample represents 0.4% of the available processor bandwidth.

17.1.4.5 Decode the Tag Data

The 168-bit raw bitstream representing the tag data has been read in. It must now be decoded into the PageID, X and Y coordinates etc. (see Table 5).

The 28×6 -bit codewords can be divided into message data codewords and checkwords. Assuming that the Version bit is 0, the 2-bit NumMsgWords field in the Mode area determines the number of message data codewords D as 15, 16, 17, or 18. The remaining K codewords are check codewords.

Each of the D 6-bit message codewords can now be examined to see if it is all 0's or all 1's. Any such codeword is illegal (due to the encoding process) and can be considered to have been erased. Determining that a codeword has been erased increases the error correction capabilities of the check bits.

The data can then be decoded using straightforward Reed-Solomon decoding. One of the advantages of using Galois Field GF(64) is that a total of only 256 bytes are required for lookup tables during the decoding process (64 bytes each of power, log, multiply and multiply inverse).

The decoding process of 168 bits to 90 bits will take approximately 4,000 cycles on a simple microprocessor. Assuming a clock speed of 25 MHz, 4,000 cycles per position sample represents 1.6% of the available processor bandwidth.

17.1.5 Build Stroke Information

The information gained from the position is added to the current stroke or is used to start a new stroke. The definition of a new stroke is that the pressure on the pen has passed a threshold to determine that the pen is now pressing on the page.

Each stroke contains the start time (in milliseconds) followed by the recovered pageID and initial position. This is followed by a series of positions for the stroke. The positions are implicitly separated by 1/100th of a second in time. While the first position is always absolute, subsequent positions are delta encoded or absolute as required. Escape codes allow compression for a number of unknown positions (the tag cannot be found), the encountering of different pageIDs (for example the initial pageID is unknown and finally a pageID is recovered from a tag during the stroke, or it may be that the pen crosses from one page to another during a stroke), and the encountering of tags with differing ContinuousPressure values, which enable or disable the inclusion of a pressure value with each position.

The following is a definition for a stroke:

```

Stroke = time PageIdRec FirstPosRec ([PosRec | PressurePos | Escape]) endStroke
PageIdRec = {unknown | pageId}
Escape = {pageId | pressureMode | TimeSkip}
TimeSkip = {unknown | zeroDelta} shortTime
PressurePos = PosRec penPressure
PosRec = {PenAbsRec | PenDeltaRec}
FirstPosRec = {unknown | PenAbsRec}
PenAbsRec = tagCoordinate imageCoordinate macrodotDelta tilt
PenDeltaRec = penDelta imageCoordinate macrodotDelta tilt

```

where:

```

time = 32 bits unsigned (units of 1000th second)
shortTime = 16 bits unsigned (units of 1000th second)
pageId = 64 bits unsigned
tagCoordinate = 2 x 10 bits unsigned (representing X and Y)
imageCoordinate = 2 x 10 bits unsigned fixed point 7:3 (representing X and Y)
macroDotDelta = 2 x 8 bits signed fixed point S3:4 (representing ΔX and ΔY)
penDelta = 2 x 4 bits signed (representing -8..7 for X and Y)
penPressure = 5 bits
tilt = 2 x 5 bits (representing 2 tilt axes)
pressureMode = 1 bit

```

The length of a sample position is 66 bits (20 + 20 + 16 + 10) when the pen coordinate is fully described, and 54 bits (20 + 8 + 16 + 10) when the pen coordinate is a delta amount. Adding flag bits to specify that there is no escape mode (1 bit), we specify 67 and 55 bits respectively (approx 7-9 bytes).

The initial position requires 165 bits to define (32 + 64 + 69). Adding 1 bit to specify that the position is known at pen-down, we require 166 bits (approx 21 bytes).

For a stroke over a single page that lasts 4 seconds, and assuming no read errors, there is an initial pen down followed by 399 deltas. The total stroke length is therefore 21,712 bits (166 + 399 × 54), which is 2,714 bytes, or 6.785 bytes per sample.

If pressure is also captured with each position (depending on the value of the encountered ContinuousPressure bits within each tag), a further 5 bits is required for each position. The Escape Code of pressureMode allows the toggling of this state within a stroke.

The stroke definition copes well with unknown pen locations. If the initial position (pageId, coordinates etc.) is unknown due to the inability to locate a tag or recover tag data, the initial position for the stroke is marked as unknown. This initial position may be followed by an Escape Code of TimeSkip, signifying the amount of time where unknown positions were captured. Finally, when a tag is captured and decoded successfully, the Escape Code for pageId is used. This pageId is valid until the next Escape Code pageId is included in the stroke (inserted when a tag is encountered with a pageId different to the current *known* pageId). If a tag cannot be found or decoded successfully mid stream within a stroke, the Escape Code for TimeSkip is used, which represents the elapsed time where only unknown positions were sampled. An unknown pen position does not change the pageId. If tag positions are reacquired, the pageId may be changed, depending on the tag. A Timeskip entry for a set of unknown positions uses 1 bit for the escape code, 2 bits to define the type of escape code, and 16 bits to define the duration of samples skipped.

The other use for the Escape Code of TimeSkip is when there have been one or more deltas of 0. The Escape Code of zeroDelta combines with a 16-bit time amount to specify how long the pen has remained in the same position without moving.

The pen contains 0.5MB of RAM to buffer stroke information both during stroke construction and if the pen is out of range from the receiver. A 0.5MB buffer allows for approximately 192 strokes at 4 seconds per stroke, which represents more than 12 minutes of stroke capture. The exact duration will depend on the success of tag decoding, and whether or not continuous pressure is being captured.

If the stroke buffer becomes full, a *Buffer-Full* LED is illuminated. The LED stays illuminated until the pen comes into range of a receiver, whereupon the buffered strokes can be transmitted and the buffer freed.

17.1.6 Activate Light

The position tag data area contains a LightEnable bit (see Table 5). The pen's HotSpot LED must be enabled directly in response to the value of this bit. There is no additional interpretation of the LightEnable bit by the pen.

17.2 DATA ENCRYPTION

The data stream must be encrypted before being transmitted to the appliance. Otherwise eavesdroppers could potentially capture signature information.

Consequently, the final stage before transmitted is the encryption of the data stream. Symmetric encryption is used since it is tractable in the time available. Possible symmetric algorithms for use are listed in Table 9. In all cases, a minimum key size of 14 bytes is assumed.

Table 9. Symmetric Algorithms [83]

Algorithm	Type	Clocks per byte processed (Pentium)	Block size (bytes)	Key size (bytes)	Patents
Blowfish	Block	18	8	4-56	No
CAST	Block	20	8	16	Yes, but free
IDEA	Block	50	8	16	Yes
Triple-DES	Block	108	8	14 ^a	No

a. Although Triple-DES uses a 168 bit key, the effective key strength is only 112 bits

Of the algorithms listed in Table 9, Triple-DES is the most conservative choice, although it is the slowest. At 108 clocks per byte processed, the encoding of a single 8-byte block consumes 864 cycles (8×108).

The pen must send 100 encrypted samples each second. Assuming that each sample requires a complete Triple-DES block (8 bytes), the time taken to encrypt 1 second of data is 86,400 cycles.

For a clock speed of 25 MHz, 86,400 cycles for the 100 samples equates to less than 0.5% of the processor bandwidth.

The encryption process therefore has a throughput of 800 bytes per second.

REFERENCES

- [1] Adobe Systems Inc., *Portable Document Format Reference Manual*, version 1.2, November 27, 1996
- [2] Adobe Systems Inc., *PostScript Language Reference Manual*, Second Edition, 1990
- [3] Bank of Montreal Economics Department, "Newsprint Market Update", Forest Products, May 1998
- [4] Bender, W., "Read all about it in the Daily You", Communicating Business, Forward Publishing, London, Winter 1994/95 (<http://nif.www.media.mit.edu/papers/forward.html>)
- [5] Bender, W., P. Chesnais, S. Elo, A. Shaw, and M. Shaw, "Enriching communities: Harbingers of news in the future", IBM Systems Journal, Vol.35, Nos.3&4, 1996, pp.369-380
- [6] Brand, S., "Personal Newspaper", in *The Media Lab*, Penguin Books, 1988, pp.36-39
- [7] Bray, T., "News Wire Services Heading for XML", 12 August 1998, <http://xml.com/xml/pub/98/08/nitf.html>
- [8] Cable Datacom News, "Cable Modem Customer Count To Top 500,000 At Year's End", December 1998, <http://www.cabledatcomnews.com/dec98-1.htm>
- [9] Cable Foreman Xtra, "WDM boosts fiber-optic capacity", May 1998, <http://www.broadband-guide.com/cbl4man/eng/eng5981.html>
- [10] Chesnais, P.R., M.J. Mucklo and J.A. Sheena, "The Fishwrap Personalized News System", *IEEE Second International Workshop on Community Networking*, Integrating Multimedia Services to the Home, 1995
- [11] Cisco Systems, "Cisco 605 Personal PCI ADSL Modem", <http://www.cisco.com/warp/public/728/605>
- [12] Cisco Systems, "Cisco 6100 Series advanced DSL Access Multiplexer", <http://www.cisco.com/warp/public/728/6100>
- [13] Clark, T., "Visa, Mastercard try to revive SET", *CNET NEWS.COM*, 12 May 1999, <http://www.news.com/News/Item/0,4,0-36435,00.html@st.ne.180.head>
- [14] Claymon, D., "When push comes to shove", *Red Herring* magazine, April 1997
- [15] Click~Through, <http://www.clickthrough.com>
- [16] CNN, *CNN Custom News*, http://customnews.cnn.com/cnews/pna_auth.welcome
- [17] CNN Interactive, "With 'electronic ink', a printed page can be a work in progress", <http://cnn.com/TECH/science/9812/23/electronic.ink.ap/index.html>
- [18] Cook, P.S., D. Gomery, and L.W. Lichty (eds.), *The Future of News - Television, Newspapers, Wire Services, Newsmagazines*, The Woodrow Wilson Center Press, 1992
- [19] Crothers, B., "IBM eyes new way to sell PCs", *CNET NEWS.COM*, 14 May 1999, <http://www.news.com/News/Item/0,4,0-36508,00.html@st.ne.ni.lh>
- [20] Cycorp, <http://www.cyc.com>
- [21] Cycorp, "Knowledge-Enhanced Searching of Captioned Information", in *Applications of CYC*, <http://www.cyc.com/applications.html>
- [22] Daniel, R., and Mealling, M., "Resolution of Uniform Resource Identifiers using the Domain Name System", *Internet Engineering Task Force RFC2168*, June 1997
- [23] Doane, W.J., "A Comparison of Display Technologies for E-Books", <http://www.jmc.kent.edu/futureprint/1998fall/doane.htm>
- [24] E Ink, <http://www.eink.com>
- [25] Ericsson, *3G Cellular Technology*, <http://www.ericsson.se/wireless/products/modsys/3rdgen/subpages/3gcell>
- [26] E&P Interactive, "Online News Reading Soars", *E&P Interactive*, 12 June 1998, <http://www.mediainfo.com/ephome/news/newshtm/stories/061298n4.htm>
- [27] E&P Interactive, "Audience Rather than Circulation", *E&P Interactive*, 24 July 1998, <http://www.mediainfo.com/ephome/news/newshtm/stories/072498n4.htm>
- [28] E&P Interactive, "Newspapers Near Death, Gates Says", *E&P Interactive*, 5 February 1999, <http://www.mediainfo.com/ephome/news/newshtm/stories/020599n1.htm>
- [29] Fishkin, K., "Filling a Region in a Frame Buffer", *Graphics Gems*, Academic Press 1990, pp.278-284
- [30] Fitzgerald, M., "Newspapers Re-think Microzoning", *E&P Interactive*, 14 January 1999, <http://www.mediainfo.com/ephome/news/newshtm/stories/011499n3.htm>
- [31] Gannet Co. Inc., *Consolidated Statements of Income 1998*, <http://www.gannet.com/annual/ar98/financials/income.htm>


- [32] Gove, A., "More than a metaphor?", *Red Herring* magazine, June 1996
- [33] Gove, A., "Judge and Journalist", *Red Herring* magazine, November 1998
- [34] Guha, R.V. and D.B. Lenat, "Enabling Agents to Work Together", *Communications of the ACM*, Vol.37, No.7, July 1994, pp.127-142
- [35] Haase, K., "FramerD: Representing knowledge in the large", *IBM Systems Journal*, Vol.35, Nos.3&4, 1996, pp.381-397
- [36] Handley, K., "All the World's a Page",
http://www.ingersoll-rand.com/compair/oct_nov_97/octnovco/pag_1.htm
- [37] Heckbert, P.S., "A Seed Fill Algorithm", *Graphics Gems*, Academic Press 1990, pp.275-277.
- [38] Hoffert, E.M. and G. Gretsches, "The Digital News System at EDUCOM", *Communications of the ACM*, Vol.34, No.4, April 1991, pp.113-116
- [39] Hughes Network Systems, "Hughes Spaceway - Wireless Broadband on Demand",
<http://www.hns.com/spaceway/spaceway.htm>
- [40] INRIA, "Light-weight Reliable Multicast Protocol",
<http://monet.inria.fr/lrmp/lrmp.html>
- [41] INRIA, "WebCanal White Paper - Global Information Broadcast",
<http://monet.inria.fr/white/index.html>
- [42] IP Multicast Initiative (IPMI), *Higher Level Protocols used with IP Multicast*, 1997,
[http://www.ipmulticast.com/...](http://www.ipmulticast.com/)
- [43] IPTC - NAA, *News Industry Text Format*, version 2.0b1, June 1998
- [44] ISO/IEC JTC1/SC29/WG1 (Joint Photographic Experts Group), *Call for contributions for JPEG 2000: Image Coding System*, <http://www.jpeg.org/public/wg1n505.pdf>
- [45] ISO/IEC JTC1/SC29/WG1 (Joint Photographic Experts Group), *JPEG 2000 requirements and profiles version 4.0*, <http://www.jpeg.org/public/wg1n1105.pdf>
- [46] Katz, J., "Online or Not, Newspapers Suck", *Wired* magazine, September 1994
- [47] Kent Displays, <http://www.kentdisplays.com>
- [48] Kelly, K. and G. Wolf, "Push!", *Wired* magazine, March 1997
- [49] Krause, J.K., "Two More 'Free' PC Firms Pin Hopes on E-commerce", *The Industry Standard*, 1 April 1999, http://www.thestandard.net/articles/article_print/0,1454,4073,00.html
- [50] Lambert, P., "Multicast Debuts on Cable Modems", *ZDNN*, 5 August 1997,
<http://www5.zdnet.com/zdnn/content/inwk/0426/inwk0006.html>
- [51] Lenat, D.B., "CYC: A Large-Scale Investment in Knowledge Infrastructure", *Communications of the ACM*, Vol.38, No.11, November 1995, pp.33-38
- [52] Liebeskind, K., "Newspapers Continue to Report Higher Revenues", *E&P Interactive*, 8 January 1998,
<http://www.mediainfo.com/ephome/news/newshrm/stories/010899n1.htm>
- [53] Loberg, C.J., "Testing the Future of All-Optical Networks", *Lightwave Xtra*, March 1998,
<http://www.broadband-guide.com/lw/reports/report3981.html>
- [54] Loizos, C., "Media: Newspaper experts flop online with New Century Network", *Red Herring* magazine, July 1998
- [55] Lyppens, H., "Reed-Solomon Error Correction", *Dr. Dobbs' Journal* Vol.22, No.1, January 1997
- [56] McCanne, S., "Overview of Reliable Multicast Protocols", *Matters Mbone Workshop*, SIGCOMM '96, August 1996, http://gaia.cs.umass.edu/sigcomm_mcast/talk1.html
- [57] McCanne, S., "SRM: Work in Progress", *Matters Mbone Workshop*, SIGCOMM '96, August 1996, http://gaia.cs.umass.edu/sigcomm_mcast/talk2.html
- [58] Macromedia, *The Flash Open File Format White Paper*,
<http://www.macromedia.com/software/flash/open/whitepaper/whitepaper.html>
- [59] Maes, P., "Agents that Reduce Work and Information Overload", *Communications of the ACM*, Vol.37, No.7, July 1994, pp.31-40,146
- [60] Marimba, *Castanet 3.2 Family Brochure*, <http://www.marimba.com/products/resource.htm>
- [61] Marimba, *Castanet 3.0 UpdateNow SDK Datasheet*, <http://www.marimba.com/products/resource.htm>
- [62] MasterCard/Visa, *SET Secure Electronic Transaction Specification*, Book 1: Business Description, Version 1.0, 31 May 1997
- [63] Medin, M., "Data Over Cable and the @Home Network",
<http://www.academ.com/nanog/feb1996/data.over.cable.html>

- [64] Miller, G.A., "WordNet: A Lexical Database for English", *Communications of the ACM*, Vol.38, No.11, November 1995, pp.39-41
- [65] MIT Media Lab, *News in the future*, <http://nif.www.media.mit.edu/nif.html>
- [66] MIT Media Lab, *Fishwrap Documentation Center*, <http://fishwrap-docs.www.media.mit.edu/docs>
- [67] MIT Media Lab, *Micromedia - Electronic Paper*, <http://www.media.mit.edu/micromedia/elecpaper.html>
- [68] Mitchell, D., "Aggregate, then segregate", *Red Herring* magazine, November 5, 1997
- [69] Mittner, G., "E-commerce sites rejoice", *Red Herring* magazine, January 4, 1999
- [70] My Yahoo, <http://my.yahoo.com>
- [71] Negroponte, N., "Less is More", in *Being digital*, Vintage Books, 1996, pp.149-159
- [72] New Media Communication, *CyberStream Fast Internet & Data Broadcasting System via Satellite*, <http://www.nmc.com>
- [73] New York Times Co., *Consolidated Statements of Income 1998*, http://www.nytc.com/pdf/consoll_1998.pdf
- [74] Neuwirth, R., "Inserts Outpace ROP", *E&P Interactive*, 24 September 1998, <http://www.mediainfo.com/ephome/news/newshtml/stories/092598n1.htm>
- [75] Neuwirth, R., "Newspaper Ad Spending on Target", *E&P Interactive*, 20 November 1998, <http://www.mediainfo.com/ephome/news/newshtml/stories/112098n2.htm>
- [76] Neuwirth, R., "Newspaper Web Losses in the Millions", *E&P Interactive*, 14 December 1998, <http://www.mediainfo.com/ephome/news/newshtml/stories/121498n1.htm>
- [77] Newspaper Association of America (NAA), *Facts About Newspapers 1998*, <http://www.naa.org/info/facts>
- [78] Noack, D., "Online News Audience Becomes More Mainstream", *E&P Interactive*, 15 January 1999, <http://www.mediainfo.com/ephome/news/newshtml/stories/011599n2.htm>
- [79] Orwant, J., "For want of a bit the user was lost: Cheap user modeling", *IBM Systems Journal*, Vol.35, Nos.3&4, 1996, pp.398-416
- [80] Platt, C., "Digital Ink", *Wired* magazine, May 1997
- [81] Riding, K., "E-Commerce: Wallet software is moving to the server. Will consumers follow?", *Red Herring* magazine, December 1998
- [82] Rorabaugh, C, *Error Coding Cookbook*, McGraw-Hill 1996
- [83] Schneier, S., "The Crypto Bomb is Ticking", *Byte Magazine*, May 1998, pp.97-102
- [84] Schwartz, E.I., "Advertising Webonomics 101", *Wired* magazine, February 1996
- [85] SET Secure Electronic Transaction, *SET Specification Book 1: Business Description*, version 1.0, May 31, 1997, http://www.setco.org/set_specifications.html
- [86] Shardanand, U. and Maes, P., "Social Information Filtering: Algorithms for Automating 'Word of Mouth'", *Proceedings of ACM Computer-Human Interaction (CHI) 1995*
- [87] Steinberg, D., "Broadening the Bandwidth to the Home in USA and Canada", 15 May 1997, <http://www.specialty.com/hiband/mi-int.html>
- [88] Sterne, J., *Advertising on the Web*, Que 1997
- [89] Stevens, W.R., *TCP/IP Illustrated, Volume 1 - The Protocols*, Addison-Wesley 1994
- [90] Stone, M., "Newspapers Must Change, Fast - API Cites Serious Threat of Online Classifieds", *E&P Interactive*, 24 July 1998, <http://www.mediainfo.com/ephome/news/newshtml/stories/072498n1.htm>
- [91] Sun Microsystems, <http://www.sun.com>
- [92] Tappert, C.C., Suen, C.Y., Wakahara, T., "The State of the Art in On-Line Handwriting Recognition", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol.12, No.8, August 1990
- [93] Teledesic, *Teledesic Overview*, <http://www.teledesic.com/overview.html>
- [94] The New York Times on the Web, <http://www.nytimes.com>
- [95] The Unicode Consortium, *The Unicode Standard, Version 2.0*, Addison Wesley 1996
- [96] Turkowski, K., "Filters for Common Resampling Tasks", *Graphics Gems*, Academic Press 1990, pp.147-165
- [97] Weinberger, D., "The Daily Me? No, the Daily Us", *Wired* magazine, April 1995


- [98] Welch Allyn Inc., *Aztec Barcode Symbology Specification Rev 3.0*, 1995
<http://dcd.welchallyn.com/techover/dcdwhite.htm>
- [99] Wicker, S., and Bhargava, V., *Reed-Solomon Codes and their Applications*, IEEE Press 1994
- [100] World Wide Web Consortium, *Cascading Style Sheets, level 2, CSS2 Specification*,
W3C Recommendation, 12 May 1998, <http://www.w3.org/TR/REC-CSS2/>
- [101] World Wide Web Consortium, *Extensible Markup Language (XML) 1.0*,
W3C Recommendation, 10 February 1998, <http://www.w3.org/TR/1998/REC-xml-19980210>
- [102] World Wide Web Consortium, *Extensible Stylesheet Language (XSL) Specification*,
W3C Working Draft, 21 April 1999, <http://www.w3.org/TR/WD-XS/>
- [103] World Wide Web Consortium, *HTML 4.0 Specification*,
W3C Recommendation, 24 April 1998, http://www.w3.org/TR/REL_html40/
- [104] World Wide Web Consortium, *Scalable Vector Graphics (SVG) Specification*,
W3C Working Draft, 12 April 1999, <http://www.w3.org/TR/WD-SVG/>
- [105] Yahoo, <http://www.yahoo.com>
- [106] Zollman, P.M., "1999 Could make or Break News Web Sites", *E&P Interactive*, 16 December 1998,
<http://www.mediainfo.com/ephome/news/newshtm/stories/121698n1.htm>

MEMJET


Micro Electro Mechanical Inkjet




0 μ s




1 μ s




2 μ s



3 μ s



4 μ s



5 μ s

Appendix B



Silverbrook Research Pty Ltd
393 Darling Street, Balmain
NSW 2041 Australia
Phone: +61 2 9818 6633
Fax: +61 2 9818 6711



MEMJET INTRODUCTION

Memjet is a new digital printing technology under development at Silverbrook Research. It is a 'clean slate' development aimed at developing the 'perfect' print technology for a wide range of applications for which current digital printing technologies are inadequate. The price/performance advantage over existing technologies such as laser printers and thermal inkjet printers is around two orders of magnitude.

Key Features of Memjet:

- Pagewidth inkjet printing - no scanning printheads, therefore very fast.
- High nozzle count: 51,200 nozzles for A4/letter.
- Full quality color photographic images at 1600 dpi.
- Full quality text, including Japanese Kanji.
- Wide ink and paper flexibility - three times less water than 600 dpi printers.
- High speed - 30 pages per minute (ppm) to 4000 ppm.
- High nozzle density - 25,600 nozzles in a 75mm² chip (compare to latest HP - 300 nozzles in 75 mm²).
- Low cost - under \$10 for a 30 ppm letter/A4 full color 1600 dpi printhead.
- Low power allows battery operation for many applications.
- Small size - printers incorporated in mobile phones, cameras, even pens.
- Simple drive circuits - 3V digital ASIC with low pincount.
- High volume manufacture - 56 million 8" heads from a 25,000 wafer per month fab.
- Low manufacturing investment - can adapt a 0.5 micron CMOS fab.
- Excellent patent protection - 220 US patents pending, both basic patents and strategic blocking patents.


Radical, not Evolutionary

The closest technology to Memjet is Thermal Ink Jet (TIJ). This technology was invented roughly simultaneously by Canon and Hewlett-Packard around 1980 (Canon's variety is known as 'Bubblejet'), and currently results in annual revenues around \$24 billion.

Thermal inkjet printheads propel a droplet of ink out of a nozzle by superheating a tiny volume of ink. This ink undergoes a flash evaporation process, forming a bubble which pushes the ink out of the nozzle.

There has been a massive investment in developing thermal inkjet technology, which has steadily advanced products from initial 200 dpi black only printers with 12 nozzles, to current full color 600 dpi printers containing printheads with as many as 608 nozzles.

However, Memjet does not use the thermal inkjet operating principle, and is not an extension of TIJ technology.



Cross section of a Memjet nozzle

Micro Electro Mechanical Systems (MEMS)

Memjet is derived from MEMS technology. MEMS is Micro Electro Mechanical Systems, and is basically the construction of mechanical systems using VLSI chip fabrication techniques. MEMS allows the integration of hundreds of millions of mechanical devices on a wafer. Certain MEMS processes (such as the Memjet process) allow the integration of MEMS and CMOS processing. This is essential for a pagewidth printhead, otherwise around 50,000 off-chip connections would be required, making the cost prohibitive for volume markets.

MEMS devices are used in many applications, though few of these applications have reached substantial volume sales. Some of the best known are

CONTENTS

Memjet Introduction	3
Advantages	6
Yield	13
Volume	15
Prototype Fabrication	17



Silverbrook Research
393 Darling Street
Balmain NSW 2041
Australia
Ph: +61 2 9818 6633
Fax: +61 2 9818 6711
info@silverbrook.com.au

accelerometers for car airbags, and Texas Instruments' DMD (Digital Micromirror Device).

The DMD is an array of around 1 million tiny mirrors on a CMOS chip. Each mirror is independently deflected in response to video image data distributed on the CMOS chip. Each mirror reflects light either towards a viewing screen, or towards a baffle. As all of the mirrors are independently controlled, a large amount of data can be projected even though the individual speed of the mirrors is only around 50 kHz.

Memjet is conceptually similar. A Memjet printhead contains around 25,000 tiny paddles on a CMOS chip. Each paddle is independently deflected in response to page image data distributed on the CMOS chip. When deflected, each paddle pushes a microscopic ink drop out of a nozzle towards the paper. As all of the paddles are independently controlled, a large amount of data can be printed even though the individual speed of the paddles is only around 40 kHz.

Memjet Value Proposition


The value proposition of Memjet is 'practical pagewidth inkjet printing'.

A scanning printhead is fundamentally around 100 times slower than a pagewidth printhead. This is because the a scanning printhead must scan across the page numerous times - around 40


times for draft mode and 200 times for photo quality.

The most fundamental difference between scanning and pagewidth printheads is the width of the printhead.

There is an essential divide between scanning printheads (roughly 1/2") and pagewidth printheads (8" for US letter/A4). Printheads between these values are not very useful - there is no 'evolutionary' progression.



*Texas Instruments
Digital Micromirror Device*



*Memjet
Inkjet nozzle*

For example: you can't make a letter/A4 pagewidth printer out of a single 4" printhead. Neither can you make a reasonable scanning printer - a 4" print swath is virtually impossible to 'stitch' on the paper without getting highly objectionable lines across the page.

A pagewidth printhead means lots of nozzles. For a 1600 dpi, 4 color, 8" printhead, 51,200 nozzles are required. This compares to typically 300 nozzles for TIJ, and 128 nozzles for piezoelectric inkjets (such as Epson and Tektronix).

Previous pagewidth printhead projects

Pagewidth printheads have long been the 'holy grail' of the inkjet industry, and have been the subject of intensive research at various large companies for more than a decade.

Previous attempts at pagewidth inkjets have mostly tried to scale up either of the two current technologies - thermal inkjet or piezoelectric inkjet. Both of these approaches fail for a number of fundamental reasons.

Why cant you make pagewidth thermal inkjet printheads?

Thermal inkjet (TIJ or Bubblejet - HP and Canon's technology) is not suitable for pagewidth printheads as the power consumption is around 100 times too high. This makes it


very difficult to get sufficient power into the printhead. However, even more difficult is getting the waste heat out of the printhead without boiling the ink.

TIJ cannot be made in a monolithic process, as the nozzle plates must be around 30 microns thick to withstand the transient pressures generated by the bubble. As a result, there are extraordinary manufacturing difficulties in scaling up from small printheads to pagewidth printheads. These include differential thermal expansion (making it extremely difficult to align the nozzles at both ends of a long printhead simultaneously) and cracking of the silicon chips due to the high stresses during nozzle plate attachment.


Also, TIJ and Bubblejet printheads wear out due to kogation and cavitation (neither of these problems affect Memjet or piezoelectric printheads). While it is cost effective to frequently replace 1 cm long scanning printheads, pagewidth printheads should have a longer lifetime.

Why cant you make pagewidth piezoelectric inkjet printheads?

The problems with scaling piezoelectric printheads are entirely different than those for TIJ. The main problem is acoustic crosstalk. As the printhead length increases, the delay between the acoustic pulse of a piezo actuator and its acoustic reflections from the ends of the printhead increase. This causes the pulse to interfere with itself on a timescale which reaches into the critical time of



Scanning printhead



Pagewidth printhead

drop ejection. Also, as the number of nozzles increases, the number of possible interference combinations undergoes a combinatorial explosion. This makes it essentially impossible to obtain consistent drop ejection at various places on the printhead.

Another problem is cost. Silicon is not a piezoelectric material, and it is extremely difficult to integrate piezoelectric materials on CMOS chips. This means that a separate external connection is required for each nozzle. A printhead equivalent to an 8" Memjet printhead would require at least 51,201 external connections. While theoretically possible, the cost of these connections is exorbitant.

The drive voltage required is around 100 V to 200 V, so it is also difficult to integrate large numbers of drivers on a single chip.

Another problem preventing integration is that the piezoelectric material must be electrically poled at around 100,000 Volts, which makes it extremely difficult to protect integrated CMOS circuitry.

How fast can a Memjet printer go?

For most consumer and PC printer applications, the full speed of Memjet is not required. In these applications, the print speed will generally be limited by the low cost paper movement mechanism.


However, for various industrial and commercial printing applications, higher print speeds are desirable.

Individual Memjet nozzles can operate at 40 kHz (and are likely to operate up to 100 kHz). At 1600 dpi, this means a print speed of 25 inches per second. A "standard" Memjet printhead has four channels for four ink colors. If this printhead is used to print a single color instead, then the maximum print speed is 100 inches per second. A 34 inch wide web (paper roll) fed printer, with printheads on both sides of the paper, printing at 100 inches per second, has an equivalent print speed of 4,364 ppm. At these print speeds, it is possible to compete directly with commercial offset printing.


Memjet - TIJ trends

The bar charts to the right show the trends for thermal inkjet printheads from the first commercial products to the present. These trends are compared to what is expected to be the first Memjet printhead - a 4", 4 color, 1600 dpi printhead with 25,600 nozzles. Two such printhead chips butted together make up an 8" pagewidth printhead.


Nozzles per monolithic printhead



Ink drops per CC



Nozzle pitch



ADVANTAGES

The list of advantages is long, and there are no disadvantages listed. This is because Silverbrook Research has spent years analyzing the problems with previous technologies, and ensuring that these problems are solved for Memjet. Memjet has been through 47 major design iterations to converge on its current state, where we believe that further improvements will not be significant. This quest for perfection is ongoing, and it is our intention to correct any undiscovered problems as quickly as possible. As is well known, a problem detected at an early stage can be orders of magnitude cheaper to correct than a problem detected at a prototype or production stage.


High Resolution

The true resolution of Memjet is 1600 dots per inch (dpi) in both directions. This allows full photographic quality color images, and high quality text (including Kanji). Higher resolutions are possible with the technology. 2400 dpi and 4800 dpi versions have been investigated for special applications, but 1600 dpi is chosen as ideal for most applications. The true resolution of commercially available advanced thermal inkjet devices is around 600 dpi. For piezoelectric systems such as those from Epson the true resolution is substantially lower, but the printhead makes many overlapping passes over each point giving an 'effective resolution' of 1440 dpi, at the expense of speed. 'Addressable resolution' is now often quoted in advertising to give the illusion of higher resolution.

Excellent Image Quality

High image quality requires high resolution and accurate placement of drops. The monolithic pagewidth nature of Memjet allows drop placement to sub-micron precision. High accuracy is also achieved by eliminating misdirected drops, electrostatic deflection, air turbulence, and eddies, and maintaining highly consistent drop volume and velocity. Image quality is also ensured by the provision of sufficient resolution to avoid requiring multiple ink densities. Five color or 6 color 'photo' inkjet systems introduce halftoning artifacts in mid tones (such as flesh tones) if the dye interaction and drop sizes are not absolutely perfect. This problem is eliminated in bi-level three color systems such as Memjet.

Nozzle firing frequency




High Speed - up to 120 ppm per printhead

The pagewidth nature of the printhead allows high-speed operation, as no scanning is required. The print speed of a Memjet based printer will usually be determined by market requirements and the paper handling mechanisms required. For most consumer or SOHO applications, 30 ppm printing is around the optimum, as faster printers are dominated by complex and expensive paper handling. To achieve 30 ppm, the Memjet printhead is operated at a drop repetition rate of 10 kHz. The maximum drop repetition rate for Memjet is 40 kHz, allowing 120 ppm full color printing from a single printhead.

Low Cost

A photo-width photographic printhead assembly is projected to cost under \$5, and a pagewidth A4 printhead assembly (using two 4" printhead equivalents) is projected to cost less than \$10 when manufactured using a 0.5 micron CMOS process on 8" wafers.



When early 300 mm fabs eventually become cost effective, monolithic 8" printheads can be fabricated on a single

wafer, reducing production costs further.

All Digital Operation

The high resolution of the printhead is chosen to allow fully digital operation using digital halftoning. This eliminates color non-linearity (a problem with continuous tone printers), and simplifies the design of drive ASICs.

Small Drop Volume

To achieve true 1,600 dpi resolution, a small drop volume is required. Memjet's drop volume is one picoliter (1 pl). The drop size of advanced commercial piezoelectric and thermal inkjet devices is around 10 to 30 pl. This has been steadily reduced over the previous two decades from original drop volumes of around 100 pl. A small drop volume also allows substantially less ink carrier (typically water) to be printed to the page, allowing much faster drying and eliminating print-through. It is theoretically possible to build a 1600 dpi thermal inkjet printhead with 1 pl drops. However, such a printhead would need to print 7.11 times as many drops as a current 600 dpi printer. It would be 7 times slower, and use around 4 times the energy, as a 600 dpi printhead with an equivalent number of nozzles. Thermal inkjet printers are already too slow and use too much power, so it is unlikely that such print-heads will be built. Memjet uses around 100 times less energy to print a drop, and uses around 100 times as many nozzles (51,200 versus around 512 for thermal inkjet), so there are no speed or power problems arising from the high resolution.

Accurate Control of Drop Velocity

As the drop ejector is a precise mechanical mechanism, and does not rely on bubble nucleation, accurate drop velocity control is available. This allows low drop velocities (4 m/s) to be used in applications where media and airflow can be controlled. Drop velocity can be accurately varied over a considerable range by varying the energy provided to the actuator. High drop velocities (10 to 15 m/s) suitable for printing on rough or fibrous surfaces (such as 'plain paper') can be achieved using variations of the nozzle chamber and actuator dimensions. Twelve meters per second is chosen as the nominal velocity for Memjet plain paper applications where airflow is uncontrolled.

Fast Drying

A combination of very high resolution, very small drops, and high dye density allows full color printing with much less water ejected. Memjet ejects around one third of the water of a 600 dpi thermal inkjet printer. This allows fast drying and virtually eliminates paper cockle.

Wide Temperature Range

Memjet is designed to cancel the effect of ambient temperature. Only the change in ink characteristics with temperature affects operation and this can be electronically compensated. Operating temperature range is expected to be 0 °C to 50 °C for water based inks.

No Special Manufacturing Equipment Required

The manufacturing process for Memjet leverages entirely from the established semiconductor manufacturing industry. Most inkjet systems encounter major difficulty and expense in moving from the laboratory to production, as high accuracy specialized manufacturing equipment is required. For Memjet, the equipment required has already been developed at a cost of many billions of dollars for the semiconductor industry.

High Production Capacity Available

An 8" CMOS fab with 25,000 wafer starts per month can produce around 56 million 8" printheads per annum. There are currently many such CMOS fabs in the world.


Low Factory Setup Cost

The factory set-up cost is low because existing 0.5 micron 6" and 8" CMOS fabs can be used. At the time of

projected introduction of the printheads these fabs should be fully amortized, and essentially obsolete for CMOS logic production. Therefore, volume production can use 'old' existing facilities. Most of the MEMS post-processing can also be performed in a CMOS fab, but at least one new piece of equipment will be required, - a deep silicon etcher using the 'Bosch process'. These machines are available from Alcatel, PlasmaTherm, and Surface Technology Systems.

Good Light- Fastness

As the ink is not heated, there are few restrictions on the types of dyes that can be used. This allows dyes to be chosen for optimum light-fastness. Some recently developed dyes from companies such as Zeneca Specialties and BASF have light-fastness of 7 on the blue wool scale. This is equal to the light-fastness of many pigments, and considerably in excess of photographic dyes and of early inkjet dyes.



Copper phthalocyanine - a blue chromophore - showing an electron density isosurface mapped with electric potential

Good Water- Fastness

As with light-fastness, the lack of thermal restrictions on the dye allows selection of dyes for characteristics such as water-fastness. For extremely high water-fastness (as is required for washable fabrics) reactive dyes can be used.

Excellent Color Gamut

The use of transparent dyes of high color purity allows a color gamut considerably wider than that of offset printing and silver halide photography. Offset printing in particular has a restricted gamut due to light scattering from the pigments used. With three-color systems (CMY) or four-color systems (CMYK) the gamut is necessarily limited to the tetrahedral volume between the color vertices. Therefore it

is important that the cyan, magenta and yellow dyes are as spectrally pure as possible. A slightly wider 'hexcone' gamut that includes pure reds, greens, and blues can be achieved using a 6 color (CMYRGB) model. Such a six-color printhead can be made economically with a width of only 1 mm.

Elimination of Color Bleed


Ink bleed between colors occurs if the different primary colors are printed while the previous color is wet. While image blurring due to ink bleed is typically insignificant at 1600 dpi, ink bleed can 'muddy' the midtones of an image. Ink bleed can be eliminated by using microemulsion-based ink, for which Memjet is highly suited. The use of microemulsion ink can also help prevent nozzle clogging and ensure long-term ink stability.

Advanced Ink Formulations

Silverbrook Research has world-class expertise in quantum chemistry (refer to papers in J. Phys. Chem. and Chemical Physics Letters), and is applying this to the development of a new class of inks with advanced properties.

High Nozzle Count

Memjet has 19,200 nozzles in a monolithic CMY three-color photographic printhead. While this is large compared to other printheads, it is a small number compared to the number of devices routinely integrated on CMOS VLSI chips in high volume production. It is also less than 3% of the number of movable mirrors which Texas Instruments integrates in its Digital Micromirror Device (DMD), manufactured using similar CMOS + MEMS processes.



Geometry of packaged photocopier printhead

51,200 Nozzles per A4 Pagewidth Printhead

A four-color (CMYK) Memjet printhead for pagewidth letter/A4 printing uses two chips. Each 0.65 cm² chip has 25,600 nozzles for a total of 51,200 nozzles.

Integration of Drive Circuits

In a printhead with as many as 100 nozzles, it is essential to integrate data distribution circuits (shift registers), data timing, and drive transistors with the nozzles. Otherwise, a minimum of 51,201 external connections would be required. This is a severe problem with piezoelectric inkjets, as drive circuits cannot be integrated on piezoelectric substrates. Integration of many millions of connections is common in CMOS VLSI chips, which are fabricated in high volume at high yield. It is the number of off-chip connections that must be limited.

Monolithic Fabrication

Memjet is made as a single monolithic CMOS chip, so no precision assembly is required. All fabrication is performed using standard CMOS VLSI and MEMS processes and materials.


In thermal inkjet and some piezoelectric inkjet systems, the assembly of nozzle plates with the printhead chip is a major cause of low yields, limited resolution, and limited size. Also, pagewidth arrays are typically constructed from multiple smaller chips. The assembly and alignment of these chips is an expensive process.

Modular, Extendable for Wide Print Widths

Long pagewidth printheads can be constructed by butting 'standard' 100 mm Memjet heads together. The edge of the Memjet printhead chip is designed to automatically align to adjacent chips. One printhead gives a photographic size printer, two gives an A4 printer, and four gives an A3 printer. Larger numbers can be used for high speed digital printing, pagewidth wide format printing, and fabric printing.

Low Cost Simultaneous Double Sided Printing

Double sided printing (known as 'duplex' in the office market, and 'perfect' in the commercial print market) can be implemented at low cost simply by including an extra printhead on the other side of the paper, and duplicating the appropriate logic and image processing circuits. The cost and complexity of providing two printheads is less than that of mechanical systems to turn over the sheet of paper.



Simultaneous high resolution printing

Straight Paper Path

As there are no drums required, a straight paper path can be used to reduce the possibility of paper jams. This is especially relevant for office duplex printers, where the complex mechanisms required to turn over the pages are a major source of paper jams.

High Efficiency

Thermal inkjet printheads are only around 0.1% efficient (electrical energy input compared to drop kinetic energy and increased surface energy). Memjet is more than 10 times as efficient.

Self-Cooling Operation

The energy required to eject each drop is 142 nJ (0.142 microJoules), a small fraction of that required for thermal inkjet printers. The low energy allows the printhead to be completely cooled by the ejected ink, with only a 32 °C worst-case ink temperature rise. No heat sinking is required.

Low Pressure

The maximum pressure generated in a Memjet printhead is around 60 kPa (0.6 atmospheres). The pressures generated by bubble nucleation and collapse in thermal inkjet and bubblejet systems are typically in excess of 10 MPa (100 atmospheres), which is 160 times the maximum Memjet pressure. The high pressures in bubblejet and thermal inkjet designs result in high mechanical stresses.

Low Power

A 30 ppm A4 Memjet printhead requires a maximum of 67 Watts when printing full 3-color black. When printing 5% coverage, average power consumption is only 3.4 Watts.

Low Voltage Operation

Memjet can operate from a single 3V supply, the same as typical drive ASICs. Thermal inkjets typically require at least 20 V, and piezoelectric inkjets often require more than 50 V. The Memjet actuator is designed for nominal operation at 2.8 volts, allowing a 0.2 V drop across the drive transistor, to achieve 3V chip operation.

Operation from 2 or 4 AA Batteries

Power consumption is low enough that a photographic Memjet printhead can operate from AA batteries. A typical 6 x 4 inch photograph requires less than 20 Joules to print (including drive transistor losses). Four AA batteries are recommended if the photo is to be printed in 2 seconds. If the print time is increased to 4 seconds, 2 AA batteries can be used.

Battery Voltage Compensation


Memjet can operate from an unregulated battery supply, to eliminate efficiency losses of a voltage regulator. This means that consistent performance must be achieved over a considerable range of supply voltages. Memjet senses the supply voltage, and adjusts actuator operation to achieve consistent drop volume.

Advanced Nozzle Clearing

Memjet employs two novel patented techniques to clear clogged nozzles, and most conventional techniques can also be used. The degree of potential nozzle clogging is highly dependent on the ink formulation, which depends on the application. In photographic applications clog-proof inks can be used. Hot melt inks can also be completely clog free. When using aqueous inks, it is essential to use nozzle capping and to include a humectant in the ink.

Small Actuator and Nozzle Area


The area required by a Memjet nozzle, actuator, and drive circuit is 3200 μm^2 . This is much less than 1% of the area required by piezoelectric inkjet nozzles, and around 1.5% of the area required by TIJ nozzles. The actuator area directly affects the printhead manufacturing cost.



Top View

Small Total Printhead Size

An entire printhead assembly (including ink supply channels) for a letter/A4, 30 ppm, 1600 dpi, four color printhead is 210 x 12 x 10 mm. The small size allows incorporation into notebook computers and miniature printers. A photograph printer is 106 x 8 x 8 mm, allowing inclusion in pocket digital cameras, palmtop PCs, mobile phone/fax, and so on. Ink supply channels take most of this volume. The photographic printhead chip itself is only 102 x 0.55 x 0.3 mm.



Printers in 3G mobile phones

Various Nozzle Capping Systems

Nozzle capping systems have been designed for various applications.


A miniature nozzle capping system has been designed for portable and photographic applications. For a photographic printer this nozzle capping system is only 106 x 5 x 4 mm, and does not require the printhead to move.

For very low cost printers, such as printers incorporated into pens, the nozzle capping mechanism is a single piece of molded plastic costing less than a cent.

In desktop printers, the printhead can be capped against a transfer roller for

robust operation and consumer-proofing.

Customized seal caps the nozzle against the transfer roller



Cross section of packaged printhead (CEprint version)

High Manufacturing Yield

The projected manufacturing yield (at maturity) of the Memjet printheads is at least 80%, as it is primarily a 0.5 micron digital CMOS chip with an area of only 0.15 cm² per inch of printhead. Most modern CMOS processes achieve high yield with chip areas in excess of 1 cm². For chips less than around 1 cm², cost is roughly proportional to chip area. Cost increases rapidly between 1 cm² and 4 cm², with chips larger than this rarely being practical. There is a strong incentive to ensure that the chip area is less than 1 cm².

For thermal inkjet and bubblejet printheads, the chip width is typically around 5 mm, limiting the cost effective chip length to 1 to 2 cm. A major target of Memjet has been to reduce the chip width as much as possible, allowing cost effective monolithic pagewidth printheads.

In the early stages of manufacture, before high yields are obtained, fault tolerance can be used. Although fault tolerance doubles the 'raw' chip area, wafers with defect densities as high as 100 defects per cm² can still obtain good printhead yields.

Low Process Complexity

With digital IC manufacture, the mask complexity of the device has little or no effect on the manufacturing cost or difficulty. Cost is proportional to the number of process steps, and the lithographic critical dimensions. Memjet uses a standard 0.5 micron 1P2M CMOS manufacturing process, with an additional 7 MEMS mask steps. This makes the manufacturing process less complex than a typical 0.25 micron CMOS logic process with 5 level metal. However, as the MEMS postprocessing


is not standard, a significant amount of process development is required. Considerable effort has been undertaken to minimize the complexity and risk of this process development. However, since any process development is usually difficult and expensive, this is likely to be the highest portion of the remaining development costs.

Simple Testing

Memjet includes test circuits and a test strategy that allows most testing to be completed at the wafer probe stage. Testing of all electrical properties, including the resistance of the actuator, can be completed at this stage. However, actuator motion can only be tested after release from the sacrificial materials. Actuators can be tested before being filled with ink using optical methods. The paddle is reflective and the nozzle deflection can be measured accurately by counting interference fringes of monochromatic light. Final testing of packaged printheads is readily performed by printing a test pattern which is automatically checked using a linear image sensor.

Low Cost Packaging

Memjet is packaged in an injection molded polycarbonate package. All connections are made using Tape Automated Bonding (TAB) technology, though wire bonding can be used as an option. All connections are along one edge of the chip.



Molded plastic package

Relaxed Critical Dimensions

The critical dimension (CD) of the Memjet CMOS drive circuitry is 0.5 microns. Advanced digital ICs such as microprocessors currently use CDs of around 0.25 microns, which is two device generations more advanced than the Memjet printhead requires. Most of the MEMS post processing steps have CDs of 1 micron or greater.

Low Stress during Manufacture

Devices cracking during manufacture are a critical problem with both thermal inkjet and piezoelectric devices. This limits the size of the printhead that it is possible to manufacture. The stresses involved in the manufacture of Memjet printheads are no greater than those required for CMOS fabrication. Memjet printheads are not sawn from the wafer, but are gently plasma etched instead.

No Scan Banding

Memjet is a full pagewidth printhead, so does not scan. This eliminates one of the most significant image quality problems of inkjet printers. Banding due to other causes (mis-directed drops, printhead misalignment) is usually a significant problem in pagewidth printheads. These causes of banding have also been addressed.

'Perfect' Nozzle Alignment

All of the nozzles within a printhead are aligned to sub-micron accuracy by the 0.5 micron stepper used for the lithography of the printhead. Nozzle alignment of two 4" printheads to make an A4 pagewidth printhead is achieved with the aid of mechanical alignment features on the printhead chips. This allows automated mechanical alignment (by simply pushing two printhead chips together) to within 1 micron. If finer alignment is required in specialized applications, 4" printheads can be aligned optically.

Controllable Drop Velocity


An accurately controlled drop velocity improves image quality, as the position of dots on the moving substrate is more accurate. An accurate drop velocity also enables a lower nominal drop velocity, which reduces power consumption. A low drop velocity requires laminar airflow, with no eddies, to achieve good drop placement on the print medium. This is achieved by the design of the printhead packaging.

For printing on 'rough' surfaces higher drop velocities are desirable. Drop velocities up to 15 m/s can be achieved using variations of the design dimensions. It is possible to manufacture printheads with a 4 m/s drop velocity, and printheads with a 15 m/s drop velocity, on the same wafer. This is because both can be made using the same process parameters.

No Misdirected Drops

Misdirected drops are eliminated by the provision of a thin rim around the

nozzle, which prevents the spread of a drop across the printhead surface in regions where the hydrophobic coating is compromised.



No Thermal Crosstalk

When adjacent actuators are energized in bubblejet or other thermal inkjet systems, the heat from one actuator spreads to others, and affects their firing characteristics. In Memjet, heat diffusing from one actuator to adjacent actuators affects both the heater layer and the bend-canceling layer equally, so has no effect on the paddle position. This virtually eliminates thermal crosstalk.

No Structural Acoustic Crosstalk


This is a major problem with piezoelectric printheads. It does not occur in Memjet.

No Fluidic Acoustic Crosstalk

Each simultaneously fired nozzle is at the end of a 300 micron long ink inlet etched through the (thinned) wafer. These ink inlets are connected to large ink channels with low fluidic resistance. This configuration virtually eliminates any effect of drop ejection from one nozzle on other nozzles.


Acoustic crosstalk between 'pods' (groups of nozzles sharing an ink channel) is almost entirely eliminated by the 300 micron long channels opening into relatively large fluid reservoirs. Acoustic crosstalk within a pod is effectively eliminated by the time delay between subsequent drop ejections within a pod. This is typically 20 microseconds, which is more than 40 times longer than the resonant frequency of the ink channel. This gives plenty of time for acous-

tic vibrations to be damped between drop ejections.



No Power Supply Crosstalk

The thick copper bus-bars in the device package, and the active power supply compensation mechanism, effectively eliminate crosstalk coupled through the power supply.



Permanent Printhead

All of the known problems that limit the life of inkjet printheads have been eliminated, allowing the printheads to be permanently installed. This dramatically lowers the production cost of consumables. Note, however, that the selling price of consumables is not necessarily related to the production cost, and need not be reduced.

No Kogation

Kogation (residues of burnt ink, solvent, and impurities, from the Japanese 'koga' for burnt rice) is a significant problem with bubblejet and other thermal inkjet printheads. Memjet does not have this problem, as the ink is not heated.

No Cavitation

Erosion caused by the violent collapse of bubbles is another problem that limits the life of bubblejet and other thermal inkjet printheads. Memjet does not have this problem because no bubbles are formed.


No Electromigration

No metals are used in Memjet actuators or nozzles, which are entirely ceramic. Therefore, there is no problem with electromigration in the actual inkjet devices. The CMOS metallization layers are designed to support the required currents without electromigration. This can be readily achieved because the current considerations arise from heater drive power, not high speed CMOS switching.

Distributed Power Connections

While the energy consumption of Memjet is fifty times less than thermal inkjet, the high print speed and low voltage results in a fairly high electrical current consumption. Worst case current for a photographic Memjet head printing in two seconds from a 3 Volt supply is 4.9 Amps. This is supplied via copper busbars to 256 bond pads along the edge of the chip. Each bond pad carries a maximum of 40 mA. On chip contacts and vias to the drive transistors carry a peak current of 1.5 mA for 1.2 μ s, and a maximum average of 12 mA.

Alternating power and signal connections on TAB film



Close-up cross section of packaged printhead

No Corrosion

The nozzle and actuator are entirely formed of glass and titanium nitride (TiN), a conductive ceramic commonly used for metallization barrier layers in CMOS devices. Both materials are at minimum chemical energy levels with respect to water, so do not corrode. Titanium nitride does not corrode or dissolve in extreme environments such as molten aluminum. It is used as the coating for the electrodes in aluminum smelters. TiN is also highly wear resistant - many watches and jewelry items are coated with TiN as it looks like gold, but has much better wear properties.

TiN does slowly oxidize in air above 500 °C, which limits the efficiency of the actuator, as the actuator efficiency is proportional to its temperature rise.

Greater efficiency can be obtained by the use of (Ti,Al)N, which has similar properties to TiN but resists oxidation up to 900 °C.

No Electrolysis


The ink is not in contact with any electrical potentials, so there is no electrolysis. This is achieved by a deliberate design feature in the actuator arm, which is a gap between the actuator loop and the paddle. This electrically isolates the paddle from the drive circuit.

No Fatigue

All actuator movement is within elastic limits, and the materials used are all ceramics, so there is no fatigue. Finite element analysis shows the maximum strain to be 0.5%.

No Friction

No moving surfaces are in contact, so there is no friction.



The actuator and paddle do not contact the nozzle chamber

No Stiction

'Stiction' is a combination of 'sticking' and 'friction', a problem common to many MEMS devices. Memjet is designed to eliminate stiction during the release of the actuators. This is achieved by the low ratio of width to thickness of the cantilever beam, in combination with the high Young's modulus of the TiN layers.

No Crack Propagation

Finite element analysis (FEA) has shown the maximum strain to be 0.5%. This is well within the crack propagation limit of the actuator, with the typical surface roughness of the TiN layers.

No Electrical Poling Required

Piezoelectric materials must be poled after they are formed into the printhead structure. This poling requires very high electrical field strengths - around 20,000 V/cm. The high voltage

requirement typically limits the size of piezoelectric printheads to around 5 cm, requiring 100,000 Volts to pole. Memjet requires no poling.

No Rectified Diffusion

Rectified diffusion - the formation of bubbles due to cyclic pressure variations - is a problem that primarily afflicts piezoelectric inkjets. Memjet is designed to prevent rectified diffusion, as the ink pressure never falls below zero.

Elimination of the Saw Street

The saw street between chips on a wafer is typically 200 microns. This would take 26% of the wafer area. Instead, plasma etching is used, requiring just 4% of the wafer area. This also eliminates breakage during sawing.

Lithography Using Standard Steppers

Although Memjet printhead chips can be as long as the wafer is wide, standard steppers (which typically have an imaging field around 20 mm square) are used. This is because the printhead is 'stitched' using identical half inch exposures. Alignment between stitches fields is not critical, as there are no electrical connections between stitch regions. One segment of each of 32 printheads is imaged with each stepper exposure, giving an 'average' of 4 printheads per exposure.

Integration of Full Color on a Single Chip

Memjet integrates all of the colors required onto a single chip. This cannot be done with pagewidth 'edge shooter' designs, such as Canon Bubblejet.

Wide Variety of Inks

Memjet does not rely on the ink properties for drop ejection. Inks can be based on water, microemulsions, oils, various alcohols, MEK, hot melt waxes, or other solvents. Memjet can be 'tuned' for inks over a wide range of viscosity and surface tension. This is one significant factor allowing a wide range of applications.


Archival Quality

With the right choice of dye and media, archival permanence significantly better than color photographs can be achieved.

Laminar Air Flow with no Eddies

The printhead packaging is designed to ensure that airflow is laminar, and to eliminate eddies. This is important, as

eddies or turbulence could degrade image quality due to the small drop size.



Printer with cartridge and airfoil

High Drop Repetition Rate can be used


The nominal drop repetition rate of a photographic Memjet is 5 kHz, resulting in a print speed of 2 seconds per photo. The nominal drop repetition rate for an A4 printhead is 10 kHz for 30+ ppm A4 printing. The maximum drop repetition rate is primarily limited by the nozzle refill rate, which is determined by nozzle chamber geometry, flow dynamics, ink pressure, and surface tension. Drop repetition rates of 40 kHz can be achieved (and 100 kHz may be achievable), allowing print speeds of 120 ppm using a single row of nozzles for each color. However, 34 ppm is entirely adequate for most low cost consumer applications.



120 ppm duplex desktop printer

For very high-speed applications, such as commercial printing, multiple printheads can be used in conjunction with fast paper handling. A suitable system for high end commercial printing can use one printhead per color per side. On a 34 inch web, this would give an effective (A4 equivalent) print speed of 10,900 ppm with a 6.25 m/s paper speed.

If the printer speed is reduced, the energy required to print a page is distributed over a longer time, so the power consumption is reduced. For operation from AA batteries this can be important.



Printer MOS circuit

as the internal resistance of the batteries limits the available power. The battery lifetime is not an important issue, as hundreds of pages can be printed from one set of batteries. The drop repetition rate can be reduced as low as desired to reduce power consumption, as the CMOS design is fully static.

Low Head-to-Paper Speed

The nominal head to paper speed of a photographic Memjet printhead is only 0.076 m/sec. For an A4 printhead it is only 0.16 m/sec., which is about a third of the typical scanning inkjet head speed. The low speed simplifies printer design and improves drop placement accuracy. However, this head-to-paper speed is enough for 34 ppm printing, due to the pagewidth printhead. Higher speeds can readily be obtained where required.

High Speed CMOS not Required

The clock speed of the printhead shift registers is only 14 MHz for an letter/A4 printhead operating at 30 ppm. For a photographic printer, the clock speed is only 3.84 MHz. This is much lower than the speed capability of the CMOS process used. This simplifies the CMOS design, and eliminates power dissipation problems when printing near-white images.

Fully Static CMOS Design

The shift registers and transfer registers are fully static designs. A static design requires 35 transistors per nozzle, compared to around 13 for a dynamic design. However, the static design has several advantages, including higher noise immunity, lower quiescent power consumption, and greater processing tolerances.

Wide Power Transistor

The width to length ratio of the power transistor is 688. This allows a 4 Ohm on-resistance, whereby the drive transistor consumes 6.7% of the actuator power when operating from 3V. This


size transistor fits beneath the actuator. Thus an adequate drive transistor, along with the associated data distribution circuits, consumes no chip area that is not already required by the actuator.

There are several ways to reduce the percentage of power consumed by the transistor: increase the drive voltage so that the required current is less, reduce the lithography to less than 0.5 micron, use BiCMOS or other high current drive technology, or increase the chip area, allowing room for drive transistors which are not underneath the actuator. However, the 6.7% consumption of the present design is considered a cost/performance optimum.

Extensive Simulation

Extensive simulations of fluid dynamic, mechanical, thermal, electrical, and other characteristics of the device have been performed. These are done using software developed at Silverbrook Research, in combination with several leading commercial software packages (Ansys, Fidap, and Matlab). Simulation is used as a 'computational microscope' which is able to 'see' microscopic stresses, temperature profiles, and fluid flow in ways impossible with physical experiments. All simulations performed at Silverbrook Research are 'causal'. That is, no assumptions are made about the motion of the ink or other aspects of the device. A simulated voltage pulse is provided to the actuator, and then we watch what happens. Simulations are performed at sufficient resolution to capture the detailed behavior of features such as satellite drops. More than 2,000 simulations have been performed, using more than 20,000 hours

of compute time on several high performance workstations.



enabled drop operation of a Memjet
chip

YIELD

Yield is, of course, a critical aspect manufacturing any chips. For this reason, we have put considerable emphasis on maximizing yield. The base CMOS process is 0.5 micron, a mature technology able to achieve very high yields on reasonable sized chips. The chip size is small: about 8 mm² per 1/2" print head segment (1.3 cm² for an 8" printhead).


The MEMS process itself is not very sensitive to particulate contamination - most feature sizes are well above 1 micron, and most layers are quite insensitive to particulate contamination.

However, it is not a good idea to be complacent about yield, especially for a new process. Accordingly, the yield estimates that we use are pessimistic. The yield is calculated from the defect

density for four different printhead configurations.

Yield ramp-up

This projection assumes a slow yield ramp-up over around 5 years. It starts with a defect density of more than 100 defects per cm², falling to around 0.2 per cm² by 2005.



Redundancy and Fault Tolerance

To allow cost-effective production of printheads in the early years when the defect density is high, redundant printheads can be used.


The business model automatically selects between four configurations of print-head, depending upon how the cost of each configuration changes with

changing yield. These configurations are:


- 1) A printhead with full redundancy (two complete rows of nozzles) made from tested and matched 1/2" printhead segments (a total of 32 half inch segments).
- 2) A printhead with full redundancy made from two rows of untested 4"

printhead segments (a total of 4 four inch segments).

- 3) A printhead with no redundancy made from 16 half inch tested segments.
- 4) A printhead with no redundancy made from 2 four inch tested segments.



Normal printhead configuration with a single set of nozzles
(51,200 total nozzles)



Redundant printhead configuration with two sets of nozzles
(102,400 total nozzles)


Printhead Yield

The overall printhead yield can be improved dramatically by the use of

redundancy / fault tolerance when the defect densities are high.

The following graph shows the 'sort yield' of 8" printheads of four different

types based on the defect density curve above.




Printhead Manufacturing Cost

The manufacturing cost of the printhead changes with yield. Although the yield of the fault tolerant configurations of printhead will always be higher than non-fault tolerant configurations, fault

tolerance is not cost effective once the defect densities are low enough. The primary reason for this is that two complete sets of nozzles are required for fault tolerance (unlike memories, where a few redundant rows and/or columns are suf-

ficient). The graph below reflects the total manufacturing costs of printheads, including device sorting and packaging. The volume of manufacture is also considered into the cost projection.



VOLUME

The manufacturing volume is predicted to go through three major phases:


1) An initial lab phase, where only around 100 wafers per month are processed, the yield are low to none, and the available number of printheads is

very limited. Defect densities above 100 per cm^2 are assumed.

2) A pilot phase, where the number of wafers processed each month is 5,000. Defect densities are assumed to be between 10 per cm^2 and 100 per cm^2 , necessitating the use of redundancy to achieve good printhead yields.

3) A full manufacturing phase, where 25,000 wafers are processed each month. Here defect densities are expected to fall from around 8 per cm^2 to 0.1 per cm^2 over the course of five years. During this phase, it is more cost effective *not* to use redundancy.

Millions of 8" Printheads per quarter for Various Configurations
Showing Optimum Configuration



Early Product Mix

The market focus during the pilot phase and the main production phase should be different. During the pilot phase, markets which are relatively insensitive to print-head cost should be targeted. These include:

- 1) Wide format printing (machine prices around \$10,000)
- 2) Network color printing (machine prices between \$3,000 to \$7,000)
- 3) Digital commercial printing (machine prices \$100,000 and up)
- 4) Photo finishing (machine prices between \$5,000 and \$20,000)

During the main manufacturing phase, high volume markets can also be targeted.

MEMJET PROTOTYPE FABRICATION

Before an integrated CMOS + MEMS prototype is made, we recommend the fabrication of a MEMS only prototype. The MEMS prototype can be made very faithfully to a full print head, with nearly identical actuator and nozzle structure. The main limitation of a MEMS only prototype is that the number of nozzles is limited, as a separate bond pad is required for each nozzle.


The prototype described here has only 15 nozzles per chip. The behavior of a few groups of 5 nozzles is a near perfect model of the entire chip performance, as the fluidic, thermal, electrical, acoustic, or mechanical coupling between 5 nozzle groups is extremely small.

A chip layout with 15 nozzles is shown below. This chip is 3 mm x 3 mm, and is replicated on a 1.2 x 1.2 cm mask set. The mask set contains 10 variants of the precise nozzle dimensions,


plus various test structures. The mask also contains a test printhead with around 3000 nozzles. As there is no CMOS on the chip, only some of the nozzles are 'wired up'. This chip is intended to test mechanical aspects of handling and packaging.

The mask has been laid out using 'Tanner Tools', and is available as Tanner or GDSII files.

The subsequent pages show the process steps, and the mask for a single nozzle unit cell.



1) 1 Micron Aluminum




PROCESS DETAILS

One micron of aluminum is deposited and etched using Mask 1. This mask includes the electrodes to the actuator, the bond pads, and the wiring between these items. It is possible to replace the aluminum with TiN wiring and bond pads. However, that would diverge further from the CMOS + MEMS design, and add process risks. The region around the nozzle chamber is on Metal1 for a 1P2M CMOS + MEMS process, while the electrodes are on metal 2.


Layer thickness	1 micron (nominal)
Thickness variation	± 50%
Linewidth	3 microns
CD Accuracy	± 0.2 microns
Alignment accuracy	Mechanical
Align to	Wafer flat
Production process	CMOS M1 and M2
Production mask	Contains wiring



Mask 1




Top View



Deposit and etch 1 micron aluminum

2) 1 Micron PECVD Nitride



PROCESS DETAILS


One micron of PECVD silicon nitride is deposited and etched using Mask 2. This mask includes the vias from the aluminum to the first TiN layer, and some relatively minor fluid control aspects. For a CMOS + MEMS process, this is the passivation layer, and will typically be 0.5 microns of glass followed by 0.5 microns of silicon nitride.

A pure nitride passivation layer is preferable, to prevent ions from the ink from diffusing through the glass.

Layer thickness	1 micron
Thickness variation	± 20%
Linewidth	1 micron
CD Accuracy	± 0.2 microns
Alignment accuracy	± 0.2 microns
Align to	Metal 1
Production process	CMOS passivation
Production mask	Same




Mask 2



Top View

Deposit and etch 1micron PECVD $\text{Si}_x\text{N}_y\text{H}_z$

3) 1.5 Microns Sacrificial Polyimide




PROCESS DETAILS


1.5 microns of spin-on photosensitive polyimide is deposited and exposed using UV light to Mask 3. The polyimide is then developed and hardbaked. 1.5 microns is the final thickness - around 3 microns of liquid is spun on, depending upon shrinkage.

The polyimide is sacrificial, so there is a wide range of alternative materials which can be used, such as glass or aluminum. Photosensitive polyimide simplifies the processing, as it eliminates deposition, etching, and resist stripping steps.


Layer thickness	1.5 microns
Thickness variation	$\pm 10\%$
Polyimide sidewalls	approx. 60°
Linewidth	1.25 micron
CD Accuracy	± 0.15 microns
Alignment accuracy	± 0.15 microns
Align to	Metal 1
Production process	Same
Production mask	Same



Mask 3




Top View



1.5 microns sacrificial photosensitive polyimide

4) Sputter 0.2 Microns TiN or (Ti,Al)N




PROCESS DETAILS

0.2 microns of magnetron sputtered titanium nitride is deposited at 300 °C and etched using Mask 4. This layer contains the actuator and part of the paddle.


In production, the resistivity of this layer of TiN should be consistent to within a few percent over the wafer.

(Ti,Al)N is preferred to TiN for high efficiency operation, as it resists oxidation at higher temperatures.


Layer thickness	0.2 microns
Thickness variation	± 5%
Linewidth	1 micron
CD Accuracy	± 0.15 microns
Alignment accuracy	± 0.1 microns
Align to	Metal 1
Production process	Same
Production mask	Same



Mask 4




Top View



0.2 microns TiN sputtered at 300 degrees C

5) 1.5 Microns Sacrificial Polyimide



PROCESS DETAILS


This step is identical to step 3.

1.5 microns of spin-on photosensitive polyimide is deposited and exposed using UV light to Mask 5. The polyimide is then developed and hardbaked. 1.5 microns is the final thickness - spin on around 3 microns depending upon shrinkage. The thickness determines the gap between the actuator and compensator TiN layers, so has an effect on the amount that the actuator bends.

Layer thickness	1.5 microns
Thickness variation	± 10%
Polyimide sidewalls	approx. 60°
Linewidth	1 micron
CD Accuracy	± 0.5 microns
Alignment accuracy	± 0.4 microns
Align to	Metal 1
Production process	Same
Production mask	Same



Mask 5




Top View



1.5 microns sacrificial photosensitive polyimide

6) 0.2 Microns Sputtered TiN




PROCESS DETAILS

This step may be identical to step 4. However, there is no efficiency advantage in the use of (Ti,Al)N, as this layer is not part of the thermal actuator. Either TiN or (Ti,Al)N may be used equivalently.


Deposit 0.2 microns of magnetron sputtered titanium nitride, at 300 °C. The TiN is etched using Mask 6.

This layer is not electrically connected, and is used purely as a mechanical component.


Layer thickness	0.2 microns
Thickness variation	± 5%
Linewidth	1 micron
CD Accuracy	± 0.15 microns
Alignment accuracy	± 0.15 microns
Align to	TiN 1
Production process	Same
Production mask	Same



Mask 6




Top View



0.2 microns TiN sputtered at 300 degrees C

7) 8 Microns Sacrificial Polyimide, Al Mask




PROCESS DETAILS



8 microns of standard polyimide is spun on and hardbaked. This thickness determines the height to the nozzle chamber roof. As long as this height is above a certain distance (determined by drop break-off characteristics), then the actual height is of little significance. As this polyimide layer is not photosensitive, it may be a filled layer to obtain a lower coefficient of thermal expansion.

A 50 nm aluminum hard mask is deposited. One micron of resist is spun on and exposed to Mask 7. The Al hardmask is etched.

Layer thickness	8 microns
Thickness variation	+ 4, - 0.5 microns
Linewidth	2 microns
CD Accuracy	± 0.5 microns
Alignment accuracy	± 0.1 microns
Align to	TiN 1
Production process	Same
Production mask	Depends on etch




Mask 7



8 microns sacrificial polyimide, with aluminum mask

8) Etch Polyimide using Oxygen Plasma




PROCESS DETAILS


Sacrificial polyimide is anisotropically plasma etched. The sidewall angle should be better than 80 degrees. The mask design shown for Mask 7 is for 90 degree sidewalls, and should be modified to suit the etch process if the etch process used varies more than ± 3 degrees from vertical.

Etch depth	11 microns
Etch stop	TiN, Si ₃ N ₄
Sidewall angle	90 degrees
Angle accuracy	± 3 degrees
Alignment accuracy	± 0.2 microns
Align to	TiN 2
Production process	Same

Uses aluminum hard mask from previous step




Top View



Etch sacrificial polyimide using oxygen plasma

9) 1 Micron Conformal Silicon Nitride



PROCESS DETAILS

1 micron of PECVD silicon nitride is deposited at 300 °C. This fills the channels formed in the previous PS polyimide layer, forming the nozzle chamber.


In the cross section, some areas appear to be large solid areas of nitride. These are actually 2 micron thick slots viewed side on.

This layer is not particularly critical. The major requirement is good adhesion to TiN. Enclosed vacuoles should not cause problems.


The nitride deposition is followed by 1 micron of polyimide, which is hardbaked.

Si ₃ N ₄ thickness	1 micron
Thickness variation	+/- 25%
Polyimide thickness	1 micron
Thickness variation	+/- 25%
Production process	Same

No Mask




Top View



1 micron conformal PECVD Si₃N₄, 1 micron polyimide


10) Etch Polyimide and Nitride




PROCESS DETAILS

The polyimide is etched down to nitride using Mask 8. The nitride is then etched down to Sac 3 polyimide using the Sac 4 polyimide as a mask.


PI Etch depth	1.5 micron
Etch variation	± 10%
Nitride Etch depth	1.5 micron
Etch variation	± 10%
Linewidth	1 micron
CD Accuracy	± 0.15 microns
Alignment accuracy	± 0.1 microns
Align to	TiN 2
Production process	Same
Production mask	Same



Mask 8




Top View



Etch of Sac 4 polyimide and nozzle roof $\text{Si}_x\text{N}_y\text{H}_z$

11) Deposit 0.25 Microns of PECVD Nitride




PROCESS DETAILS

0.25 microns of conformal PECVD silicon nitride is deposited at 300°C.

This layer forms the nozzle rims, using a 'sidewall spacer' like process. The thickness is not particularly critical, and could be substantially thinner if desired, as there is insignificant fluidic pressure acting on the rim.

Si ₃ N ₄ thickness	nominal 0.25 microns
Thickness variation	± 25%
Production process	Same

No Mask




Top View



Deposit 0.25 microns of PECVD Si_xN_yH_z

12) Anisotropic Etch of Nitride




PROCESS DETAILS


The nozzle rim nitride is anisotropically plasma etched. The etch can be timed, as etch depth is not critical. Substantial overetch is required to ensure that only vertical nitride walls remain, and that nitride over sloping topography is completely removed.

Etch depth	0.5 microns
Depth variation	$\pm 20\%$
Production process	Same
Production mask	None

No Mask




Top View



0.5 micron anisotropic 'sidewall' etch of $\text{Si}_x\text{N}_y\text{H}_z$

13) 4 Microns of Softbaked Resist




PROCESS DETAILS

Spin on 4 microns of resist and softbake.


This resist layer is to protect the front side of the wafer during backetch. The resist thickness is to cover the topography of the MEMS devices, and thereby allow a vacuum chuck to be used.

Resist thickness	4 microns nominal
Thickness variation	+50% - 25%
Production process	Same
Production mask	None

No Mask




Top View



4 microns softbaked resist as a protective layer

14) Back-etch using Bosch Process




PROCESS DETAILS

The wafer is thinned to 300 microns (to reduce back-etch time), and 3 microns of resist on the back-side of the wafer is exposed to Mask 10. Alignment is to metal 1 on the front side of the wafer. This alignment can be achieved using an IR microscope attachment to the wafer aligner. The wafer is then placed on a platter and etched to a depth of 330 microns (allowing 10% overetch) using the deep silicon etch "Bosch process". This process is available on plasma etchers from Alcatel, PlasmaTherm, and Surface Technology Systems.


Etch depth	330 micron
Etch variation	$\pm 10\%$
Linewidth	40 microns
CD Accuracy	± 2 microns
Alignment accuracy	± 2 microns
Align to	Metal 1
Production process	Same
Production mask	Full wafer



Mask 9 (includes chip edges)




Bottom View, rotated around horizontal axis



Back-etch through wafer using Bosch process

15) Strip all Sacrificial Material




PROCESS DETAILS


The chips were diced by the previous Bosch process back-etch. However, the wafer is still held together by 11 microns of polyimide. The wafers must now be turned over. This can be done by placing a tray over the wafer on the platter, and turning the whole assembly (platter, wafer and tray) over while maintaining light pressure. The platter is then removed, and the wafer (still in the tray) is placed in the oxygen plasma chamber.

All of the sacrificial polyimide is ashed in an oxygen plasma.

No Mask




Top View



Package, bond, prime, and test

16) Package, Bond, and Prime.





PROCESS DETAILS


Glue the chip into a package with an ink inlet hole, for example, a pressure transducer package. The ink hose should include a 0.5 micron absolute filter to prevent contamination of the nozzles.

The prototype Memjet chips are 3 mm square, but the ink inlet hole region is only about 240 x 160 microns, in the center of the chip. Glue the chip into the package so that the chip ink inlet is over the hole in the package. This requires only 500 micron accuracy. Wire bond the 6 connections to nozzles to be tested.

Fill the packaged printhead under approx. 5 kPa ink pressure to prime it.



17) Test




PROCESS DETAILS


Testing of the inkjets is performed by a mixture of physical observation and simulation.

The primary physical test is to take high speed photomicrographs of the ink drops in flight. This is done with a special optical microscope, a high speed flash, a gated photomultiplier plate, and a high speed CCD camera. The images are stored on computer and carefully compared to the simulation results. When the simulation matches the experimental images to within a few percent, then the simulation can act as a 'microscope' able to see almost any aspect of the nozzle operation.


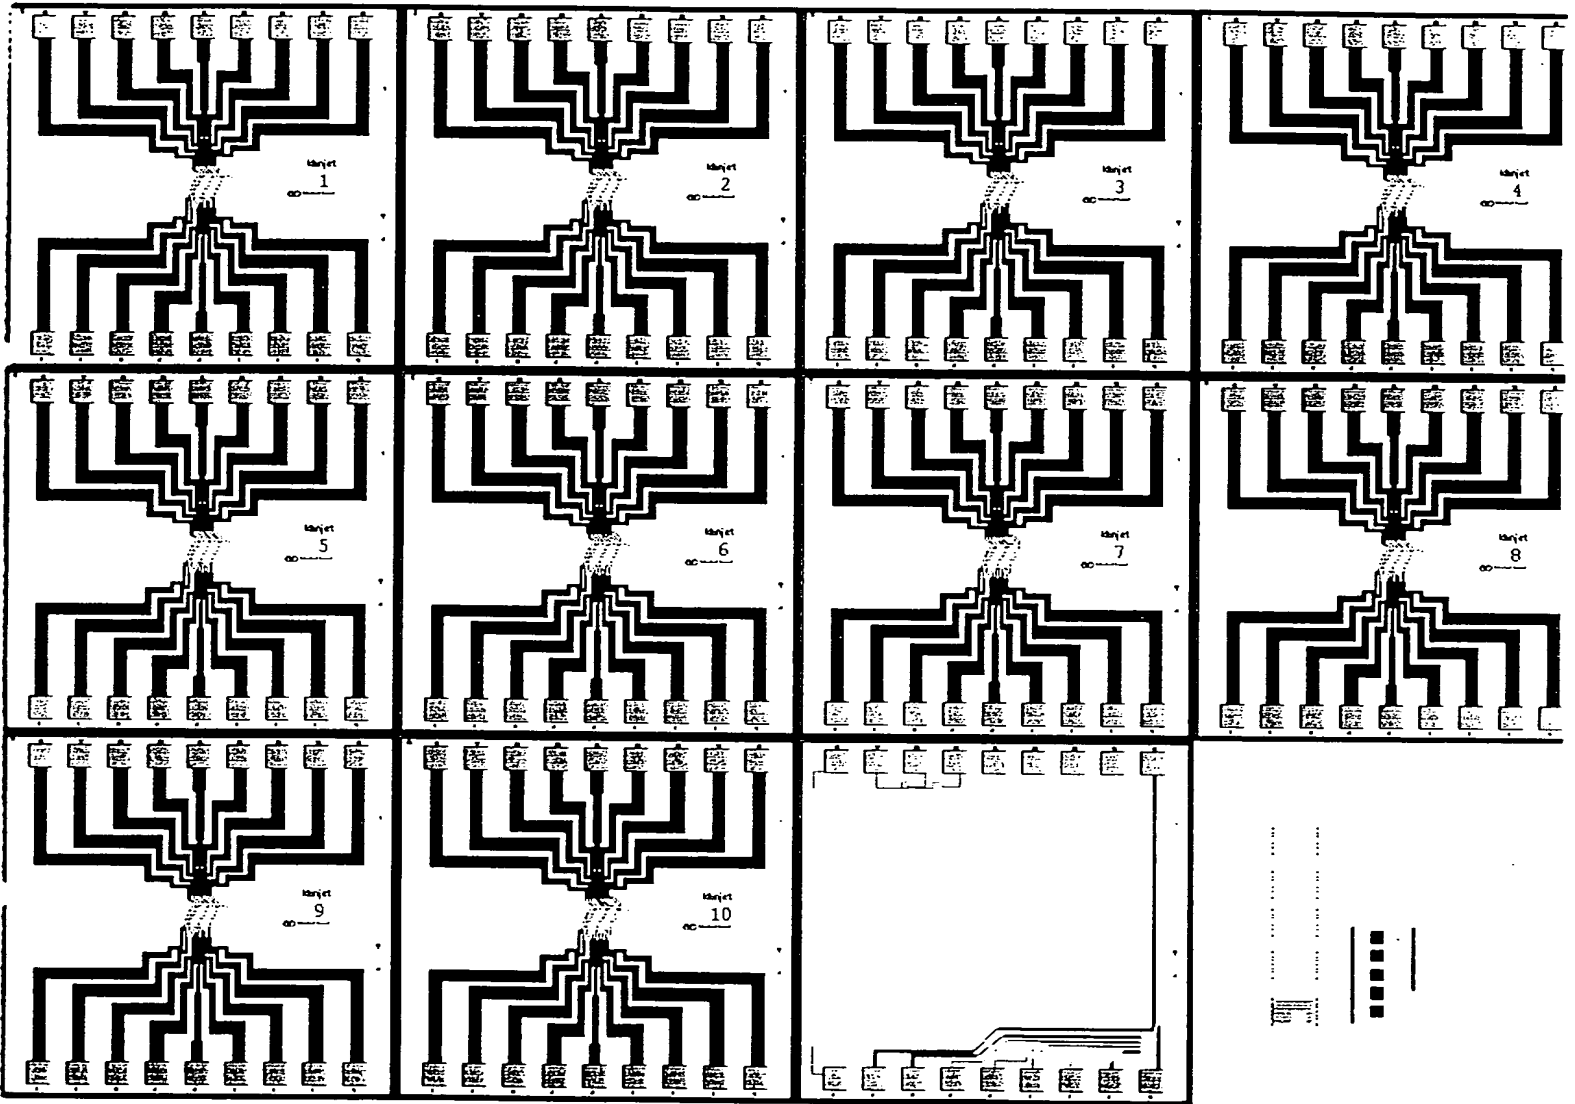
Silverbrook Research is establishing a testing laboratory that will be able to perform the necessary tests, and correlate these tests with extensive simulation results.



Actuate



Return




Appendix C

MONDAY, NOVEMBER 23, 2001

No. 465

The New Age of Robot Pets



"Twinky"

The world's first self-aware home entertainment robot was launched today by Sony. Twinky uses the latest in Artificial Intelligence, with 256 bit parallel processor and a 600,000 pixel CMOS image sensor. See story page 2.



The Mysterious Return of Tutankhamen



The mask of Tutankhamen, safely returned to the Cairo museum.

CAIRO -- As mysteriously as it disappeared two weeks ago, the famous golden mask of King Tutankhamen has been returned to its home. Egyptian police were alerted to by telephone of a package sitting outside the back entrance to the Cairo museum yesterday morning. Upon investigating further, the police discovered to their joy that it was, in fact, the missing artifact.

The artifact was discovered missing by a security guard in the museum at 4am on November the 11th. The security crackdown has led to the unearthing of hundreds of illicitly stolen artifacts throughout Egypt and the world.

Pieces that have been missing for generations have been recovered from private collections around the world, as "dob a robber" has become the favourite pass time of the rich

and famous privvy to these illicit collections. The Egyptian government has been overwhelmed with offers of help, and a crack anti-theft squad was instantly dispatched by the French police in charge of the prestigious Louvre museum in Paris.

A breakthrough came last Thursday with the discovery of the artifact's pedestal in a Cairo warehouse in the north of the city. Police were in the process of questioning several suspects when the golden death mask was returned yesterday.

"We can't believe our luck" said Mohommed Sharim, Director of the Cairo Museum. "Never before has such an audacious theft occurred, despite extremely sophisticated electronic monitoring equipment. We are very pleased the mask has been returned to us."

The New Age of Electronic Pets

TOYKO-- The latest "must have" toy has arrived on the supermarket shelves. It's called "Twinky" and it's a fully functioning robot. Twinky was designed by the Sony corporation as an intelligent pet, ready for the home, without the maintenance required for a regular animal.

Twinky is a complex design, crafted for real-world interaction. Twinky pools the talents of engineers in software, computers, sensors and com-

munications. In fact Twinky is the culmination of years of research into artificial intelligence and robotics.

Twinky uses the latest in CPU technology, and redefines the term "robot". Twinky, you see, has emotions. She can be happy or sad, bad or boisterous. And the complex sensors encapsulated into Twinky's infrastructure means that Twinky can actually see and hear the world around her.


Twinky is intelligent and autonomous, and is able to perform requested tasks and is also able to learn. She can be naughty though, and discipline is part of the robot's appeal. No spanking for Twinky, just strong words will suffice.

And don't worry about having to pick Twinky up when she falls down, she can get up all by herself.

Sales of Twinky start today, and it is anticipated that the queues at department stores will quickly be replaced by happy customers.



Sam's Email



Hi Sam!

Hi Sam,
This is Christine's Mum. Christine says to say hello and would you come over tonight and play? Sam says she hasn't been at school last week because she's been sick,

but she misses you and would like to see you soon. She has to go to the doctor's today, but she is very much better. She drew you this picture of your dog while she was sick.

Sincerely

Caroline Brown

Reply



A Ron Howard film. STARRING Cynthia Rothrock, George Touris, Barbara Littlefair
PRODUCED by Stella Berwoan, A DELTAWORKS PRODUCTION

Ick!

Sam,
I went over to Toby's house yesterday, and he was digging a cubby house in his backyard. He threw worms at me! Horrible boy! Anyway I guess we should get him back on Saturday at Chris's Party, what do you think?

Suzi Woods

Reply

Hate and other stuff

Sam, I really hate Toby Withers. Let's put dirt in his bag today. Either that or we can pull his hair. What do you think?

Shelly Wordrow

Reply

Did u know?

Hi ya Sam,
I mised you after our vacation together, but it was god fun wasn't it? Did u know that my dad is going to New York this week on a bisnes trip? He's staying in a big hotel in the city. Wot's it like living ni New York? Here in Maryland it's really quite but its gud fun on the farm.

My Mom says that I can use her computer to send you messages once a week if you would like, and we are going to Coney Island again next year for our vacation to.

Are you going again? That would be reeally cool, cause we culd play in the sand again like we did this year.

This email is taking me so long to write cause I dont know where any of the keys are. You hav e to learn to type to use a computer really fast. So when I'm older I'm going to. Bye now.

Katherine May

Reply

Compose



Gazing at the Night Sky

Clue: What is the name given to stargazing and what is it?

M	A	R	S	A	P	O	L	A	R	E	T	T	A	N
O	S	A	K	U	Q	U	A	S	A	R	E	F	M	R
O	T	Y	Y	S	N	E	L	S	T	S	T	A	K	U
N	R	I	O	L	A	H	R	S	O	T	A	R	R	T
C	O	M	E	T	L	E	O	N	A	A	T	C	A	A
E	N	E	B	U	L	A	S	T	E	R	O	I	D	S
L	O	T	O	N	I	V	E	N	U	S	R	S	M	T
E	M	E	Y	I	G	E	M	I	T	T	B	U	S	A
S	Y	O	R	V	H	N	I	G	H	T	I	P	R	R
T	O	R	O	E	T	H	L	E	S	E	T	E	E	B
I	I	A	G	R	D	S	K	C	P	R	I	R	T	U
A	D	M	R	S	E	D	Y	Y	A	R	X	N	I	R
L	A	M	I	E	N	A	W	S	C	A	N	O	P	S
G	R	A	V	I	T	Y	A	A	E	C	E	V	U	T
O	F	G	A	L	A	X	Y	S	R	T	A	A	J	R
C	O	S	M	O	L	O	G	Y	S	F	I	E	L	D

WORDS

Emit	Jupiter	Night	Saturn	USA
Asteroids	Leo	Orbit	Scan	Venus
Astronomy	Lens	Pulsar	Sky	Virgo
Celestial	Light	Polar	Star	Xray
Comet	Mars	Quasar	Starburst	
Cosmology	Meteor	Radio	Space	
Craft	Milky Way	Ray	Supernova	
Dark Matter	Moon	Rotate	Sun	
Day	Nebula	Red Dwarf	Terra	

Help

Solution

More

The words may be backwards, frontwards, diagonal or around corners.

DID YOU KNOW?

If you're having trouble remembering the names of all the planets in the solar system, simply remember this catch phrase: "My

very earnest mother just sat upon the north pole."


Each letter at the beginning of each word stands for a planet.

My = Mercury

very = Venus

earnest = Earth
just = Jupiter
sat = Saturn
upon = Uranus
north = Neptune
pole = Pluto.





Tickets

Sessions

Knock Knocks

More


Knock knock.
Who's there.
Pencil.
Pencil who?
Pencil fall down if you don't
have a belt.

Knock Knock.
Who's there?
Ice cream soda.
Ice cream soda who?
Ice cream soda neighbours
wake up!
(I scream so the neighbours
wake up!)



Tickets

Sessions



BOY SOAKER

It's here! The amazing squeezable Boy soaker!
Simply fill with water and spray those pesky
boys out of your life!
The boy soaker comes in an amazing range of
colors for the ultimate accessory!

\$9.95

Also available in
all major toy
stores



Family and Friends



Chris's Birthday Party

Hey Everybody! Don't forget it's Chris's third birthday party on Saturday! Here's a picture that I took of her last week, cool hey! Anyway, the party starts at 12, bring a plate and wine. Could you RSVP please, so I can do a head count on how many children are coming, we've got some special treats planned.

So be there, or be square.

5/315 114th St.

New York

Ph 55568484

Love Karen.

Post

Reply

Death in the Family

Grandma tells me that Great Uncle Warren passed away peacefully in the Golden Years Nursing home in NYC yesterday. Warren was a private in World War II, when he ran way to the army at the ripe old age of 17. He rose to the rank of Sergeant during the 1950's, when her served in Korea. He never married. After he left the forces he studied at NYU and then set up his own dental surgery in Mystic, Connecticut. As an active member of the local municipal council, helped to fund the Mystic Seaport.

He lost contact with many of the family in the 1960's when he moved to Africa, where he worked for Community Aid Abroad for 15 years. There he worked further in dentistry, where he was involved in training many of the younger generation of dentists in Zimbabwe and Zaire. He returned to the US in 1980, where he set up his practice in Brooklyn. There he was known in the local community for his passion for breeding butterflies. Warren Sheldon was 83.

Lucy Sheldon.



Jonathon does us Proud!

Jonathon, who has been struggling in the junior league baseball, has made the B grade in the under 12's baseball team next season!

He's playing for the Junior Braves, and hopefully he will get up to A-reserve soon. Even tho the baseball season is well and truly over, the selectors had a special playoff two weeks ago to pick the teams.

The B team will be travelling to Dallas, Texas mid season next year for two games against the B grade down there!

I took this picture on our new digital camera last summer.

Hope to hear from you all soon now!

Sally Wright

Reply

Reply



Good Luck Charm


Oh let today be sunny, inside my heart as well as outside in the world. Let me succeed in my goals and conquer my fears.

I will not let others get me down, and I will follow my dreams. For dreams are all that we can achieve, all that we could wish for.

Tickets

Sessions


Textures




[More](#) [Preview](#)
[Print](#) [Extras](#)

Coloring In




Draw lines from the surrounding colors, textures or fills to the region to be colored in. Print when finished.



Colors




Fills




The Interactive Adventures of Prince Zelmar

Yesterday's Story



As the "Tiger of Rome" sails westward towards the straits of Aroth, Crown Prince Zelmar ponders his future.



As the warships gather in the south, he now knows that hopes of peace are futile.



Princess Stephanie has rejected him, his hopes of settlement through marriage will not succeed.




Lurking deep beneath the "Tiger of Rome" lies a danger more perilous than his many enemies.






What the???

Suddenly the water beside the "Tiger of Rome" boils.




Crown Prince Zelmar gasps in astonishment as a Zoithawian sea monster burst out of the sea.



The beast is enormous! Quite capable of eating the ship as a single mouthful. But Zoithawian sea monsters are intelligent, so should Zelmar talk or fight? Continued tomorrow...



Cats of Asia



The Leopard, is an elegant, long-bodied cat with a proportionally smaller head, sturdy legs of medium length and a long tail, more sinuous in its movements than the heavier bodied lion and tiger.

Leopards are found in a wide range of habitat, from desert to savanna and forest. Once found through most of Africa and through the Middle East to Java, they are still fairly common in some parts of central and east Africa and there are scattered populations in east Africa and many parts of Asia. The close soft coat of the leopard is short and sleek in tropical areas, longer in colder climates.

In India it often seems to prefer to live near villages where it preys on dogs, goats, calves and chickens. The leopard can be a very secretive animal and is a stealthy hunter. Depredations on domestic livestock may be outweighed by their control of baboons, cave rats, and other animals that destroy crops. Other prey includes almost anything from antelopes to rodents, hares, frogs, and even dung beetles. They are strong swimmers and climb well, often carrying carcasses 20 feet (6 m) up into a tree.

The leopard can be a very secretive animal and is a stealthy hunter. In areas where it is hunted, the Leopard is nocturnal and seldom seen except in reserves where it has become used to vehicles and may be seen hunting in the daylight.

Although solitary, there have been reports of males staying with their mates and even helping rear cubs, and group of up to six have occasionally been seen.

 **Print**

The common leopard is found in parts of Asia.


 **Print**

The black Panther is a black leopard which has black rosettes on a black background.

Leopards have up to four cubs in their litters after a 98 day gestation period.

 **Print**



The Snow Leopard is an endangered species that live in mountainous regions in Asia and the Former Soviet Union.



Print

Up to four cubs (but usually two) are born after a gestation period of about 98 days, in a cave, hollow tree or thicket. The cubs are born with dull gray hair. They are weaned at three months, but stay with their mother until 18 months to two years old and reach sexual maturity and full size at two and a half to three years old.

She may move them from place to place. When the cubs are older, she will lead them to a kill. The young play by stalking and pouncing, preparing them to hunt on their own. Life span of the leopard is 21 years in captivity.

Leopards have a call, described as a "grunt-ha! grunt-ha!", and a variety of other grunts, coughs, and a sound like someone sawing wood, but they are usually silent.

A few leopards appear to be black because they have black rosettes on a black back grounds. These leopards are known as black panthers.

Leopards have been heavily hunted for their fur, and loss of habitat and natural prey has increased predation on stock (and sometimes attacks on people) leading to further killings. They are now protected through most of their range and are listed as endangered.

SNOW LEOPARD (Panthera uncia)

The snow leopard has a thick, beautiful gray to yellowish-gray coat marked with large black rosettes. Hunting for its beautiful coat is just one of the reasons this cat has become an endangered species. In some areas, natural prey sources have been depleted due to human encroachment.

Snow leopards live in mountainous regions, between 6,000 and 20,000 feet in altitude, and are found in the mountainous regions in the former Soviet Union, Mongolia, India,

Pakistan, Afghanistan, China, Bhutan and Nepal.

Although accurate wild population estimates are difficult to make due to the remoteness of the snow leopard's range, only 5,000 animals are believed to survive in the wild.

There are currently over 600 snow leopards in captivity worldwide. More than 250 animals are managed in North America as part of the American Zoo and Aquarium Association's SSP. (Species Survival Plans) The plans outline breeding plans, work to increase public awareness and education, help conduct research, and in some cases, organise programs to reintroduce captive-bred wildlife into secure habitats.

The snow leopard's range is spread across eight international boundaries, making the establishment of a cooperative effort to save them very difficult.



More



To celebrate the new season's snowfall, today's digital designs are "Creatures of Winter". You can download each of these designs at an A4 size.

This page is sponsored by *Wildlife International*. *Wildlife International* is an organisation actively involved in conservation worldwide, including the sponsorship of the "Tiger Reserve" in India, as well as purchasing large tracts of wilderness and creating managed wildlife reserves in South America. Projects also include the White Rhino breeding project in Africa and Australia, and the Bald Eagle Project here in the U.S.A. Simply press the image you want and then the Print button.

Print

More



More Stuff

Jokes

[Lightbulb Jokes](#)
[Children's Jokes](#)
[More Knock Knocks](#)
[General Jokes](#)
[Long Jokes](#)
[Blonde Jokes](#)
[Irish Jokes](#)

Join the Dots

Picture puzzles

[More Spot the Difference](#)
[Pythagorian Mindbenders](#)

Riddles

Maths Puzzles

Cross Word Puzzles

[Easy Cross Word](#)
[Cryptic Cross Word](#)
[Giant Cross Word](#)

Tongue Twisters

Comics

[Xara-Queen of the Dead](#)
[Batman](#)
[Superman](#)
[Ninja Turtles](#)
[Micky Mouse](#)
[Donald Duck](#)
[The Avenger](#)
[Invisible Man](#)
[Danger Duck!](#)
[Denis the Menace](#)
[Popeye the Sailor](#)

Art

[Landscape Art](#)
[European](#)
[African](#)
[Australian](#)
[Asian](#)
[American](#)
[The Tropics](#)

Animals

[Apes](#)
[Butterflies](#)
[Cats](#)
[Dogs](#)
[Elephants](#)
[Frogs](#)
[Giraffes](#)
[Horses](#)

Iguanas

[Jaguars](#)
[Kangaroos](#)
[Lions](#)
[Monkeys](#)
[Nautilus](#)
[Octopus](#)
[Porpoise](#)
[Rabbits](#)
[Snakes](#)
[Tortoises](#)
[Uranguatan](#)
[Vipers](#)
[Yaks](#)
[Zebras](#)

General Art

[Still Life](#)
[Portraits](#)

Collectables

[Figurines](#)
[Barbie dolls](#)
[Dolls of the World](#)
[Teapots](#)
[Doll House Furniture](#)
[Comics](#)
[Action Men](#)
[Transformers](#)
[Star Wars Figures](#)

Nintendo games

General Classifieds

[Aquaculture - Fish](#)
[Computers](#)
[Clubs and Societies](#)
[For Sale](#)
[Garage Sales](#)
[Gardening](#)
[Handicrafts](#)
[Hobbies](#)
[Internet Services](#)
[Lost and Found](#)
[Musical Instruments](#)
[Music CD's](#)
[Nintendo Games](#)
[Pets](#)
[Personal Messages](#)
[Pools](#)
[Toys](#)
[Sporting Equipment](#)

Newsfiles Menu

Do you know what you're missing out on? Want to refine your Newsfiles? Press Here.

Menu




YOUR HARDWARE HANDYMAN SPECIALISTS

Appendix D



HAND TOOLS

CONTINUED




Info



Hacksaw

With the durability of stainless steel this hacksaw will last for years

Retail: \$8.00
Our price \$6.50




Info



Pruning Saw

Tackle those pruning needs with this durable and versatile saw

Retail: \$2.00
Our price \$1.50



Info



Plaster Saw

A must for the home renovator. Trim your plaster to the exact size.

Retail: \$15.00
Our price \$11.50




Info



Carpenter's Saw

This saw has been designed to stay sharp longer.

Retail: \$18.00
Our price \$11.50




Info



Pliers (Set)

SIDCUT pliers offer you quality without the large price tag

Retail: \$9.00
Our price \$6.00




Info



Grips (Set)

SIDCUT Grips offer you quality without the large price tag.

Retail: \$4.00
Our price \$1.50




Info



Needlenosed Pliers

For those hard to get at positions you need Piontant Pliers

Retail: \$2.00
Our price \$1.50



Info




Adjustable Grips

This set of handy grips allow you to grip almost anything.

Retail: \$6.00
Our price \$4.50

POWER TOOLS




Info



HITACHI Cordless Drill

You never have to worry about a power outlet with this quality product

Retail: \$58.00
Our price \$46.50




Info



Electric Plane

Save time an effort with this professional electric plane.

Retail: \$118.00
Our price \$106.50




Info



Bench Grinder

This grinder has all the power you'll ever need.

Retail: \$125.00
Our price \$ 96.50



Info



Jigsaw

With this jigsaw you can cut just about anything cleanly


Retail: \$38.00
Our price \$29.50



**EQUIP
HARDWARE**

YOUR HARDWARE HANDYMAN SPECIALISTS

INFORMATION



HITACHI CORDLESS DRILL

TECHNICAL INFORMATION:		OPTIONAL EXTRAS:
WEIGHT:	1 POUND	Hard Carry Case
CHUCK SIZE:	1/2 INCH	Soft Carry Case
RPM:	1500 RPM	Belt Hook
CHARGE LIFE:	2 Hours Constant Use	Easy Out Attachment
RECHARGE TIME:	12 Hours Trickle Charge 1 1/2 Hour Quick charge	
BATTERIES:	2 x 9.6v NCA High Charge	












Info

Press Info Button
for Technical Manual



SHOPPING TROLLEY

QTY	DESCRIPTION		EACH	TAX	TOTAL
1	HITACHI CORDLESS DRILL 9.6v 1/2 INCH CHUCK		46.50	0	46.50
1	Sidchrome Socket Set Imperial 1/2 inch		52.50	0	52.50
1	Electric Plane Makita Power Plane		106.50	0	106.50
1	Bench Drill Ryobi with stand 4 inch chuck		180.00	0	180.00
3	Screwdriver set Sidchrome Philips Head		2.00	0	6.00
1	Pliers Bullnosed Mechman Electricians Pliers		6.00	0	6.00
1	Needlenosed Pliers Piontant Pliers Drop Forged		1.50	0	1.50
1	Adjustable Grips Sidchrome 1/2 inch and 2 inch		4.50	0	4.50
1	Carpenters Saw 14 inch Blade		11.50	0	11.50

Customer

Mr I.M. Busy
1 Apple Lane
Cambridge, MA 02139

Transaction Number: 23367
Date Printed 11/3/2001

 SUBTOTAL **234.00**









 FREIGHT **4.00**

 TOTAL **238.00**


**EQUIP
HARDWARE**

YOUR HARDWARE HANDYMAN SPECIALISTS

CHECKOUT

QTY	DESCRIPTION		EACH	TAX	TOTAL
1	HITACHI CORDLESS DRILL 9.6v 1/2 INCH CHUCK		46.50	0	46.50
1	Sidchrome Socket Set Imperial 1/2 inch		52.50	0	52.50
1	Electric Plane Makita Power Plane		106.50	0	106.50
3	Screwdriver set Sidchrome Philips Head		2.00	0	6.00
1	Pliers Bullnosed Mechman Electricians Pliers		6.00	0	6.00
1	Needlenosed Pliers Piontant Pliers Drop Forged		1.50	0	1.50
1	Adjustable Grips Sidchrome 1/2 inch and 2 inch		4.50	0	4.50
1	Carpenters Saw 14 inch Blade		11.50	0	11.50

Customer

Mr I.M. Busy
1 Apple Lane
Cambridge, MA 02139

Transaction Number: 23367**Date Printed 11/3/2001****SUBTOTAL** **235.00****FREIGHT** **4.00****TOTAL** **239.00****Payment Details**









Credit Card

☐ MASTER CARD 4564 xxxx xxxx 9696☐ Other Credit Card☐ Cash On Delivery☐ Account**10% Off****Catalog****Purchase**



YOUR HARDWARE HANDYMAN SPECIALISTS

PICKING LIST

	DESCRIPTION	QTY	Picked
	HITACHI CORDLESS DRILL 9.6v 1/2 INCH CHUCK	1	✓
	Sidchrome Socket Set Imperial 1/2 inch	1	✓
	Electric Plane Makita Power Plane	1	✗
	Screwdriver set Sidchrome Philips Head	3	✓
	Pliers Bullnosed Mechman Electricians Pliers	1	✓
	Needlenosed Pliers Piontant Pliers Drop Forged	1	✓
	Adjustable Grips Sidchrome 1/2 inch and 2 inch	1	✓
	Carpenters Saw 14 inch Blade	1	✓

Customer

Mr I.M. Busy
1 Apple Lane
Cambridge, MA 02139

Transaction Number: 23367

Date Picking List Printed 11/3/2001





YOUR HARDWARE HANDYMAN SPECIALISTS









Customer

Mr I.M. Busy
1 Apple Lane
Cambridge, MA 02139

Transaction Number: 23367

Date 11/15/2001

RECEIPT

QTY	DESCRIPTION		EACH	TAX	TOTAL
1	HITACHI CORDLESS DRILL 9.6v 1/2 INCH CHUCK		46.50	0	46.50
1	Sidchrome Socket Set Imperial 1/2 inch		52.50	0	52.50
0 backordered	Electric Plane Makita Power Plane		106.50	0	106.50
3	Screwdriver set Sidchrome Philips Head		2.00	0	6.00
1	Pliers Bullnosed Mechman Electricians Pliers		6.00	0	6.00
1	Needlenosed Pliers Piontant Pliers Drop Forged		1.50	0	1.50
1	Adjustable Grips Sidchrome 1/2 inch and 2 inch		4.50	0	4.50
1	Carpenters Saw 14 inch Blade		11.50	0	11.50

SUBTOTAL	128.50
FREIGHT	4.00
TOTAL	132.50

Equip would like to thank you for your order.

We look forward to having you shop with us again.



Appendix E

Email

from Joe Citizen 11.20pm, Wednesday, Dec 5, 2001

TO :

CC :

SUBJECT:



 **Contacts**

 **Help**

 **Attach**

 **Print**

 **Send**



Email from Joe Citizen 11.20pm, Wednesday, Dec 5, 2001

TO : *george@memjet.com*

CC :

SUBJECT: *re:cats*

Hey George check out this article on the big cats I found lying around.

For now I am preparing for a trip to the Antarctic ... but I'll be sure to catch up with you before I go.

Take care

Joe



Contacts

Help

Attach

Print

Send



Cats of Asia



The Leopard, is an elegant, long-bodied cat with a proportionally smaller head, sturdy legs of medium length and a long tail, more sinuous in its movements than the heavier bodied lion and tiger.

Leopards are found in a wide range of habitat, from desert to savanna and forest. Once found through most of Africa and through the Middle East to Java, they are still fairly common in some parts of central and east Africa and there are scattered populations in east Africa and many parts of Asia. The close soft coat of the leopard is short and sleek in tropical areas, longer in colder climates.

In India it often seems to prefer to live near villages where it preys on dogs, goats, calves and chickens. The leopard can be a very secretive animal and is a stealthy hunter. Depredations on domestic livestock may be outweighed by their control of baboons, cave rats, and other animals that destroy crops. Other prey includes almost anything from antelopes to rodents, hares, frogs, and even dung beetles. They are strong swimmers and climb well, often carrying carcasses 20 feet (6 m) up into a tree.

The leopard can be a very secretive animal and is a stealthy hunter. In areas where it is hunted, the Leopard is nocturnal and seldom seen except in reserves where it has become used to vehicles and may be seen hunting in the daylight.

Although solitary, there have been reports of males staying with their mates and even helping rear cubs, and group of up to six have occasionally been seen.

Print

The common leopard is found in parts of Asia.



Print

The black Panther is a black leopard which has black rosettes on a black background.

Leopards have up to four cubs in their litters after a 98 day gestation period.

Print

Email from Joe Citizen 11.20pm, Wednesday, Dec 5, 2001

TO : george@memjet.com

CC :

SUBJECT: re: cats

Hey George check out this article on the big cats I found lying around.

For now I am preparing for a trip to the Antarctic ... but I'll be sure to catch up with you before I go.

Take care

Joe Citizen

The Leopard, is an elegant, long-bodied cat with a proportionally smaller head, sturdy legs of medium length and a long tail, more sinuous in its movements than the heavier bodied lion and tiger.

Leopards are found in a wide range of habitat, from desert to savanna and forest. Once found through most of Africa and through the Middle East to Java, they are still fairly common in some parts of central and east Africa and there are scattered populations in east Africa and many parts of Asia. The close soft coat of the leopard is short and sleek in tropical areas, longer in colder climates.

In India it often seems to prefer to live near villages where it preys on dogs, goats, calves and chickens. The leopard can be a very secretive animal and is a stealthy hunter. Depredations on domestic livestock may be outweighed by their control of baboons, cave rats, and other animals that destroy crops.



Other prey includes almost anything from antelopes to rodents, hares, frogs, and even dung beetles. They are strong swimmers and climb well, often carrying carcasses 20 feet (6 m) up into a tree.

The leopard can be a very secretive animal and is a stealthy hunter. In areas where it is hunted, the Leopard is nocturnal and seldom seen except in reserves where it has become used to vehicles and may be seen hunting in the daylight.

Although solitary, there have been reports of males staying with their mates and even helping rear cubs, and group of up to six have occasionally

Contacts

Help

Attach

Print

Send



To: George <george@memjet.com>

From: Joe Citizen<j.citizen@citizenplace.com>

Date: Wednesday, December 5, 2001 11.20pm

Subject: re: cats

Hey George check out this article on the big cats I found lying around.

For now I am preparing for a trip to the Antarctic ... but I'll be sure to catch up with you before I go.

Take care

Joe

The Leopard, is an elegant, long-bodied cat with a proportionally smaller head, sturdy legs of medium length and a long tail, more sinuous in its movements than the heavier bodied lion and tiger.

Leopards are found in a wide range of habitat, from desert to savanna and forest. Once found through most of Africa and through the Middle East to Java, they are still fairly common in some parts of central and east Africa and there are scattered populations in east Africa and many parts of Asia. The close soft coat of the leopard is short and sleek in tropical areas, longer in colder climates.

In India it often seems to prefer to live near villages where it preys on dogs, goats, calves and chickens. The leopard can be a very secretive animal and is a stealthy hunter. Depredations on domestic livestock may be outweighed by their control of baboons, cave rats, and other animals that destroy crops. Other prey includes almost anything from antelopes to rodents, hares, frogs, and even dung beetles.



They are strong swimmers and climb well, often carrying carcasses 20 feet (6 m) up into a tree. Depredations on domestic livestock may be outweighed by their control of baboons, cave rats, and other animals that destroy crops. Other prey includes almost anything from antelopes to rodents, hares, frogs, and even dung beetles. They are strong swimmers and climb well, often carrying carcasses 20 feet (6 m) up into a tree.

The leopard can be a very secretive animal and is a stealthy hunter. In areas where it is hunted, the Leopard is nocturnal and seldom seen except in reserves where it has become used to vehicles and may be seen hunting in the daylight.

Although solitary, there have been reports of males staying with their mates and even helping rear cubs, and group of up to six have occasionally been seen.

